université PARIS-SACLAY

Advanced methods for laser control and optimisation of laser-plasma accelerators

Méthodes avancées de contrôle des lasers et d'optimisation des accélérateurs laser-plasma

Thèse de doctorat de l'université Paris-Saclay

École doctorale n° 576,Particules, Hadrons, Énergie et Noyau : Instrumentation, Imagerie, Cosmos et Simulation (PHENIICS)
Spécialité de Doctorat: Physique des accélérateurs
Graduate School :Physique. Référent : Faculté des sciences d'Orsay

Thèse préparée dans les unités de recherche **Laboratoire de Physique des 2 Infinis Irène Joliot-Curie, IJCLAB UMR9012**, sous la direction de **Fabian ZOMER** Professeure des universités, la co-direction de **Kevin CASSOU** Ingénieur de recherche.
Thèse soutenue à Paris-Saclay, le 10 Juin 2026, par

Gueladio KANE

Composition du jury

Membres du jury avec voix délibérative

JACQUES ROBERT Pr., Laboratoire de physique des gaz et des plasmas (LPGP)	Président
Arie IRMAN Dr., Helmholtz-Zentrum Dresden-Rossendorf (HZDR)	Rapporteur et examinateur
Maxence THEVENET Dr., Deutsches Elektronen-Synchrotron (DESY)	Rapporteur et examinateur
Stéphane SEBBAN Dr., Laboratoire d'optique appliquée (LOA)	Examinateur
Maud BAYLAC Dr., Laboratoire de Physique Subatomique et de Cosmologie (LPSC)	Examinatrice

THESE DE DOCTORAT
NNT : 2026UPASP033

université PARIS-SACLAY

ÉCOLE DOCTORALE
Particules, hadrons, énergie et noyau : instrumentation, imagerie, cosmos et simulation (PHENIICS)

Titre: Méthodes avancées de contrôle des lasers et d'optimisation des accélérateurs laser-plasma
Mots clés: Laser,Plasma,Acceleration,Diagnostics,Apprentissage automatique

Résumé: Cette thèse porte sur le développement de méthodes d'apprentissage automatique et d'outils de diagnostic pour un meilleur contrôle des accélérateurs laser-plasma. Dans un premier temps, l'obtention des premiers faisceaux d'électrons produits par l'injecteur laser-plasma PALLAS situé à IJCLAB est présentée. Les travaux se sont ensuite concentrés sur le développement de modèles substitutions numériques entraînés à partir de simulations Particle-In-Cell (PIC). Ces modèles permettent d'effectuer des prédictions rapides et constituent des outils utiles pour guider la recherche de points de fonctionnement optimaux. La thèse présente également l'utilisation d'une caméra multispectrale comme outil de diagnostic en intensité et en phase du faisceau laser. Une étude numérique quantifiant l'impact des déformations thermiques du compresseur optique sur la qualité du faisceau laser est également présentée. Enfin, grâce à l'utilisation de ressources de calcul haute performance, des simulations PIC réalistes de l'accélérateur ont été réalisées.

Title: Advanced methods for laser control and optimisation of laser-plasma accelerators.
Keywords: Laser,Plasma,Acceleration,Diagnostics,Machine learning

Abstract: This thesis focused on the development of machine learning methods and diagnostic tools for improved control of laser–plasma accelerator technology. First, the production of the first electron beam from the PALLAS laser–plasma injector at IJCLAB is presented. The work then focused on the development of surrogate models trained on particle-in-cell simulations. These models enable fast predictions and are useful for guiding the search for optimal operating points. The thesis also presents the use of a multispectral camera as an intensity and phase diagnostic tool for the laser beam. In addition, simulations quantifying the impact of thermal deformation of the optical compressor on laser beam quality are discussed. Finally, using high-performance computing resources, realistic particle-in-cell simulations of the accelerator were performed

Résumé étendu en français

L'accélération laser-plasma (ALP) est une méthode permettant d'accélérer des électrons et qui repose sur la génération de champs électriques longitudinaux très importants, de l'ordre de 100 GV/m. Ces champs électriques sont produits par la génération d'une onde plasma électronique, elle-même résultant de l'interaction d'un laser intense et temporellement ultracourt, de l'ordre de la femtoseconde, avec une puissance allant du TW au PW, avec un gaz. Le laser intense a deux rôles principaux au cours de cette interaction. Il va à la fois ioniser les atomes du gaz, ce qui va créer un plasma, et il est également à l'origine d'une perturbation de la densité du plasma via la force pondéromotrice, qui est à l'origine de l'onde plasma électronique.

Le laser intense est typiquement produit par la technique d'amplification par dérive de fréquence (CPA). La technique CPA consiste à étirer temporellement l'impulsion laser, qui va ensuite être amplifiée puis finalement recompressée temporellement à l'aide d'un compresseur optique. Cette technique est nécessaire afin d'éviter la génération d'effets non linéaires dans l'impulsion laser, ainsi que l'endommagement des composants optiques de la chaîne CPA. L'accélération laser-plasma est une méthode prometteuse d'accélération qui a comme principal intérêt d'être capable d'accélérer des électrons à de hautes énergies, de l'ordre de la centaine de MeV jusqu'à une dizaine de GeV, dans des structures ultracompactes de l'ordre du mm au cm. Malgré de nombreuses améliorations au cours des deux dernières décennies quant à la génération d'électrons par ALP, de nombreux défis restent à relever pour rendre cette technologie viable. En effet, l'ALP est toujours en retard par rapport aux méthodes traditionnelles d'accélération en ce qui concerne la qualité et la stabilité des électrons produits. Pour améliorer le contrôle des paramètres laser et plasma et, en définitive, la qualité et la stabilité des électrons produits, plusieurs voies de recherche sont actuellement poursuivies par de nombreux laboratoires. On peut citer notamment l'intégration progressive de méthodes d'apprentissage automatique permettant un contrôle plus fin de ces paramètres. Également, le développement de techniques de diagnostic du laser permettant une caractérisation complète du laser en monocoup, avec pour objectif final la mise en place de techniques de contrôle fonctionnant également en monocoup.

La plateforme PALLAS, située à l'IJCLab, s'inscrit dans ce cadre de recherche et a pour objectif la production de faisceaux d'électrons présentant une stabilité et une qualité comparables à celles obtenues avec les méthodes traditionnelles d'accélération, à savoir une énergie médiane comprise entre 150 et 250 MeV, une dispersion en énergie inférieure à 5 %, une charge supérieure à 30 pC et une émittance transverse inférieure à 2 mm.mrad. Cette plateforme utilise le laser produit par LASERIX, qui fournit un faisceau fonctionnant à 10 Hz, avec une puissance maximale de 62,5 TW, pour une énergie de 2,5 J et une durée de 40 fs.

Le premier chapitre est consacré aux rappels théoriques nécessaires à la compréhension de l'ALP. Sont notamment introduits les principaux concepts d'aberrations optiques ainsi que les couplages spatio-temporels du faisceau laser, effets importants car ils dégradent la qualité du faisceau laser et impactent, par conséquent, les propriétés du faisceau d'électrons produit.

Le deuxième chapitre décrit la plateforme PALLAS, justifie les choix de conception de la ligne accélératrice et détaille les différents outils de diagnostic du laser, du plasma et des faisceaux d'électrons. Il rapporte également les premiers faisceaux d'électrons obtenus sur la plateforme.

Le troisième chapitre porte sur l'élaboration de modèles de substitution numériques fondés sur l'apprentissage automatique. Ces modèles sont entraînés à partir d'un grand nombre (plus de 3000) simulations Particle-In-Cell (PIC) basse fidélité. L'intérêt de tels modèles est qu'ils offrent la possibilité de prédire les propriétés des faisceaux d'électrons en un temps négligeable par rapport aux simulations PIC. Ils peuvent ainsi être exploités pour une exploration rapide de l'espace des paramètres et servir d'outil d'aide à la conception de futures lignes ALP. Au cours de ce chapitre, nous démontrons que ces méthodes sont capables d'atteindre un haut niveau de précision dans leurs prédictions. Nous utilisons ensuite le meilleur modèle pour prédire des points de fonctionnement optimaux de la plateforme PALLAS. Enfin, une approche similaire est mise en œuvre pour prédire les spectres des faisceaux d'électrons.

Le quatrième chapitre est consacré à une caméra multispectrale destinée à mesurer l'intensité spatio-spectrale du faisceau laser en monocoup. Tout d'abord, la calibration de cette caméra, réal-

isée au Center for Advanced Laser Applications (CALA) à Munich, est décrite. Ensuite, en s'appuyant sur l'algorithme de Gerchberg–Saxton, cet outil de diagnostic est exploité afin de reconstruire la phase spatio-spectrale du faisceau laser à partir de données simulées. Enfin, une approche alternative fondée sur une architecture de deep learning permettant d'obtenir la phase du faisceau laser est proposée.

Le cinquième chapitre s'intéresse à l'une des principales sources de couplages spatio-temporels dans les chaînes CPA : le compresseur optique. Dans un premier temps, un modèle permettant de calculer la modification de l'angle de diffraction induite par la déformation du compresseur est développé. Les déformations du compresseur de LASERIX sont modélisées à l'aide du solveur COMSOL Multiphysics. En s'appuyant sur ce modèle ainsi que sur les déformations calculées avec COMSOL, un code de calcul de longueur de trajet optique est développé. Ce code permet finalement de représenter qualitativement l'impact de ces déformations sur les propriétés du faisceau laser au point focal. Par ailleurs, une campagne expérimentale utilisant la caméra multispectrale et le faisceau laser fonctionnant à 10 Hz est menée. Le chapitre s'achève par une comparaison entre les résultats expérimentaux et les simulations.

Le dernier chapitre est consacré à l'analyse de simulations PIC réalistes utilisant une décomposition en modes azimutaux. Trois configurations partageant les mêmes paramètres plasma sont étudiées : un cas basse fidélité, un cas considérant le laser fonctionnant à 1 Hz et un cas utilisant le laser fonctionnant à 10 Hz. Ces simulations mettent en évidence le développement de couplages spatio-temporels pour un faisceau laser initialement dépourvu de tels effets. Enfin, un modèle préliminaire mettant en évidence un effet de *pulse shaping* du faisceau laser est présenté en annexe.

Remerciements

Tout d'abord, j'aimerais te remercier, Fabian, pour m'avoir accepté en tant qu'étudiant. J'aimerais également te remercier pour les discussions au sujet de l'effet SISSPM en fin de thèse ainsi que pour tes encouragements. Je t'en suis très reconnaissant.

Merci beaucoup à toi, Sophie, pour tout le temps que nous avons passé ensemble à revoir chacun des chapitres de la thèse ainsi que l'article sur le surrogate model. Merci également pour les nombreux Zoom que nous avons faits ensemble et pour tes remarques pertinentes qui m'ont permis de mettre un point final à ce manuscrit.

Merci Kevin, sans qui, vraiment, je pense que je n'aurais pas pu terminer cette thèse. Merci pour la patience dont tu as fait preuve durant ces quatre années. Je te suis vraiment reconnaissant pour le soutien indéfectible que tu m'as apporté malgré les difficultés de santé auxquelles tu as dû faire face, que ce soit pour la relecture des différents manuscrits, les nombreuses discussions scientifiques ou les encouragements à poursuivre mes différentes idées tout au long de cette thèse. Tu es un super directeur de thèse et, encore une fois, merci.

Merci aux membres du jury d'avoir accepté de relire ce long manuscrit et d'avoir participé à ma soutenance.

Merci Moana pour l'aide apportée au cours des différentes manips que j'ai effectuées durant cette thèse. En particulier, merci pour les longues soirées où nous sommes restés tard au laboratoire afin de réaliser les mesures sur l'impact du chauffage du compresseur sur le faisceau laser.

Merci Pierre pour l'aide que tu m'as apportée dans la compréhension des simulations PIC en début de thèse. Merci également pour le temps que tu as pris pour relire l'article *SM* ainsi que pour les notebooks que tu m'as fournis et qui m'ont permis de gagner énormément de temps.

Merci Hayg et Hadrien pour les incroyables cours de boxe que nous avons suivis ensemble et qui m'ont permis de souffler au cours de ce doctorat. Hayg, tu es un incroyable professeur de boxe. Merci pour les techniques à la Roy Jones Jr.

Je tiens à te remercier Olivier Neveu pour ton aide dans mon introduction au code Python, notamment avec l'utilisation des classes, et en particulier pour ton aide dans le développement des différentes interfaces graphiques.

Merci également à Marceau, qui a réussi à débloquer un bug lors du développement de l'interface graphique de la caméra multispectrale.

Mykyta, merci pour les nombreuses discussions sur le machine learning que nous avons eues. J'aimerais également te remercier pour l'aide que tu m'as apportée avec les simulations réalistes lors de la réponse aux reviewers. Enfin, merci pour les nombreuses pauses café durant lesquelles nous avons bien rigolé.

Jana, merci pour les notebooks et les simulations CFD sur lesquels je me suis appuyé pour écrire mon chapitre consacré aux simulations réalistes. Surtout, merci pour la bonne humeur que tu as apportée au laboratoire.

Slava Kubytsky, je te remercie pour ton aide lors de mon introduction aux différentes méthodes de machine learning. Je te remercie également pour les fous rires que nous avons eus ensemble lors de l'EAAC.

Slava Smartsev, je tiens à te remercier même si nous n'avons pas eu le temps de travailler ensemble. Tu es quelqu'un de très chaleureux et merci pour tes remarques pertinentes lors de la préparation de ma soutenance.

Je tiens à te remercier, Bruno. Tu étais un super professeur à Polytech. C'est notamment grâce à tes cours sur les lasers que j'ai décidé de poursuivre en thèse. Merci également pour ton aide au cours des premières expériences sur la cellule de gaz que nous avons menées durant mon stage de Master 2.

Merci à Olivier Guillbaud pour les notes de cours que tu m'as fournies et qui m'ont permis de préparer les travaux dirigés que j'ai assurés à Polytech avec les APP4. Cela m'a été d'une très grande aide et je t'en suis très reconnaissant.

Merci à Kevin Dupraz et à Serge Carrez pour l'opportunité que vous m'avez donnée de dispenser des cours à Polytech. Ce fut une expérience très enrichissante.

Merci à Kalyani pour les nombreuses discussions que nous avons eues au sujet des simulations

réalistes Particle-In-Cell. Surtout, je tiens à te remercier pour tout le temps que j'ai passé sur tes comptes du TGCC et du CINES. Sans toi, je n'aurais pas pu terminer mon dernier chapitre.

Merci Nicolas d'avoir accepté de faire partie de mon comité de suivi individuel. Merci également pour les plats que tu as concoctés avec Sophie, c'était délicieux. Enfin, merci pour l'opportunité que j'ai eue lors de mon stage de M1 de faire partie de ton équipe. Ce stage a confirmé l'intérêt que je porte à la physique.

Merci à Elsa pour le temps que nous avons passé à travailler sur le laser. J'ai beaucoup apprécié travailler avec toi, tu es quelqu'un de très gentil.

Merci à toi, Marie-Hélène. Même si nous n'avons pas beaucoup interagi ensemble, tu as toujours été positive lorsque nous nous rencontrions et cela m'a fait chaud au cœur.

Merci à Francesco pour ton expertise sur les simulations PIC ainsi que pour le temps que tu as consacré à la relecture de mon article. Cela m'a permis de le finaliser.

Merci Arnaud pour les discussions autour des simulations PIC ainsi que pour les codes que tu nous as envoyés concernant la décomposition en modes azimutaux.

Enfin, je tiens à remercier ma famille : ma petite sœur Hawa, qui m'apporte de la joie au quotidien ; mon père Yero, qui me pousse à devenir quelqu'un de meilleur ; et ma mère Houleye, qui a été présente dans toutes les épreuves de ma vie. Je n'aurais pas pu rêver d'une meilleure famille.

Enfin, je tiens à remercier ma grande sœur Raki, que j'ai perdue durant ces années de doctorat. Tu étais quelqu'un de formidable. Je pense à toi très souvent et tu me manques beaucoup. Que ton âme repose en paix.

A ma grande sœur Raki

Contents

Introduction **13**

1 Laser wakefield theory **19**
1.1 Chirped pulse amplification . . . 19
1.1.1 Oscillator . . . 19
1.1.2 Stretcher and compressor . . . 25
1.1.3 Amplification . . . 29
1.2 Gaussian beam formalism . . . 32
1.3 Optical aberrations . . . 34
1.4 Spatio-temporal/spectral couplings . . . 37
1.4.1 Theory of first-order STC . . . 38
1.5 Plasma parameters . . . 40
1.5.1 Electronic density . . . 41
1.5.2 Plasma frequency . . . 41
1.5.3 Critical density . . . 41
1.6 Ionisation processes . . . 42
1.7 Electrons in an electromagnetic wave . . . 43
1.8 Fluid description of laser plasma interaction . . . 45
1.8.1 Linear solution . . . 47
1.8.2 1D Nonlinear plasma waves . . . 53
1.9 Particle trapping . . . 54
1.10 Blowout regime . . . 57
1.10.1 Injection mechanisms . . . 59
1.10.2 Self-focusing . . . 62
1.10.3 self-compression . . . 63
1.10.4 Limits of laser–plasma acceleration . . . 63
1.11 Electron parameters . . . 64
1.11.1 Electron spectrum . . . 65
1.11.2 Emittance . . . 66
1.11.3 Twiss parameters . . . 66
1.11.4 Beam loading . . . 67

2 PALLAS a Laser-plasma injector test facility **73**
2.1 Introduction . . . 73
2.2 Objectives . . . 73
2.3 LASERIX - laser driver . . . 73
2.3.1 LASERIX frontend . . . 74
2.3.2 LASERIX amplification stages . . . 74

2.3.3 Laser beam transport line . . . 75
2.3.4 Optical compressor . . . 76
2.4 LPI accelerator beamline . . . 76
2.5 Laser injection module . . . 76
2.6 Plasma module . . . 78
2.6.1 Double target cell directly integrated in the beamline . . . 78
2.6.2 Gas injection system . . . 78
2.6.3 Plasma diagnostics . . . 81
2.7 Electron characterization beamline . . . 81
2.7.1 Laser removal . . . 81
2.7.2 Electron beam diagnostics . . . 82
2.8 First results from PALLAS electron source commissioning . . . 89
2.8.1 Laser properties . . . 90
2.8.2 Electron beam properties . . . 91

3 Surrogate Models studies for laser-plasma accelerator 101
3.1 Motivation . . . 101
3.2 Machine learning methods . . . 102
3.2.1 MLP . . . 102
3.2.2 Gaussian process regression . . . 105
3.2.3 Decision trees . . . 107
3.2.4 Model evaluation . . . 108
3.3 Numerical experiment . . . 110
3.4 Data sets generation . . . 110
3.4.1 Injection prediction model . . . 113
3.5 Model construction . . . 116
3.5.1 Training surrogate model construction resources requirement . . . 117
3.6 Results . . . 118
3.6.1 Performances of surrogate models . . . 118
3.6.2 Optimisation with the surrogate models . . . 124
3.6.3 Comparison with random scan optimisation . . . 126
3.7 Spectrum prediction . . . 130
3.8 Conclusion and perspective . . . 133
3.8.1 Output electron beam parameters evaluation . . . 134

4 Spatio-Spectral Characterization and Phase Retrieval 139
4.1 Introduction . . . 139
4.2 Multi spectral camera . . . 139
4.3 Calibration . . . 139
4.3.1 Calibration setup . . . 140
4.3.2 Image processing . . . 140
4.4 Multispectral camera sensitivity . . . 143
4.5 Phase retrieval . . . 145

4.5.1 Angular spectrum decomposition . . . 146
4.5.2 Gerchberg–Saxton algorithm . . . 147
4.5.3 Phase retrieval on simulated data . . . 148
4.6 Discussion . . . 152
4.7 Algorithmic development for phase retrieval . . . 154
4.7.1 Dataset generation . . . 154
4.7.2 Machine learning architecture . . . 155
4.7.3 Loss function . . . 158
4.7.4 Conclusion . . . 162

5 Compressor deformation **165**
5.1 Introduction . . . 165
5.2 MGE derivation . . . 166
5.3 Code structure . . . 168
5.3.1 Deformation calculation Using COMSOL MULTIPHYSICS . . . 168
5.3.2 Deformed compressor code . . . 173
5.3.3 Ray propagation at deformed gratings . . . 174
5.3.4 Near-field definition and far-field propagation . . . 178
5.4 Results . . . 180
5.4.1 1 Hz laser operation . . . 180
5.4.2 10 Hz laser . . . 185
5.4.3 Simulation results and comparison . . . 189
5.4.4 Conclusion . . . 194

Appendix A: Effect of grating pitch variations **197**

Appendix B: Solution of the modified grating equation (MGE) **201**

6 Realistic PIC simulation for LWFA **205**
6.1 Introduction . . . 205
6.2 PIC theory . . . 206
6.3 Gaussian temporal envelope simulation . . . 208
6.3.1 Gaussian temporal envelope and plasma parameters . . . 208
6.3.2 Gaussian temporal envelope simulation results . . . 209
6.4 1 Hz Laser simulation . . . 213
6.4.1 1 Hz Laser parameters . . . 213
6.4.2 1 Hz Laser simulation results . . . 217
6.5 10 Hz Laser simulation . . . 224
6.5.1 10 Hz Laser parameters . . . 224
6.5.2 10 Hz Laser simulation results . . . 227
6.6 Conclusion . . . 232

Appendix C: Vacuum simulation **233**

Appendix D: Self interference spatio spectral phase modulation **237**

Appendix E: Full expression of the group delay derivative **253**

Conclusion **257**

Introduction

The main question I have received from close friends and family throughout my PhD is some variation of: What is the purpose of your work? What are laser-plasma accelerators useful for? This introduction seeks to provide a clear answer.

There are many types of particle accelerators; however, they all share a common goal: to bring particles, whether leptons (electrons, for example) or hadrons (protons, neutrons, etc.), to high energies. These high-energy particles can then be used in a wide range of scientific and technological applications.

One of the first modern particle accelerator was designed in the 1930s by Cockcroft and Walton [1] to accelerate protons and study nuclear disintegration processes. In the following decades, particle accelerators took many different forms, including linear accelerators (LINACs), cyclotrons, and synchrotrons [2].

These accelerators have found applications across a remarkably broad range of disciplines. In fundamental physics, they allow us to investigate the properties of matter at the smallest scales. One of the most notable recent discovery was the Higgs boson at CERN in 2012 [3]. In medicine, high-energy protons or X-rays are used to target and destroy tumoral cells. In industry, accelerators are employed in the fabrication and engraving of integrated circuits. Even fields such as archaeology benefit from accelerator-based techniques, for example through radiocarbon dating using accelerator mass spectrometry. These examples illustrate just a few of the many applications of particle accelerators across science and technology.

All of the previously mentioned types of accelerators share one important limitation: they rely on large metallic structures and extensive facilities to accelerate particles. This is because the longitudinal electric fields in these accelerators, typically generated using radio-frequency (RF) cavities, cannot exceed approximately 100 MV/m without causing electrical breakdown. As a result, achieving high particle energies requires very long structures, such as the 27 km circumference of CERN's Large Hadron Collider.

To overcome this limitation and enable acceleration in much more compact structures, plasma-based acceleration was proposed by Tajima and Dawson in 1979 [4]. The advantage of using plasma as an accelerating medium is that it can sustain much higher longitudinal electric fields, on the order of 100 GV/m. Exploiting such high accelerating gradients would make it possible to reach high particle energies within structures ranging from millimeters to centimeters in length, producing electron beams from the MeV to GeV scale. Tajima and Dawson showed that a short and intense laser pulse could drive a plasma wave, called a wakefield, capable of accelerating electrons to high energies.

The first attempts to generate plasma waves using lasers in the 1980s and early 1990s relied on the so-called beatwave acceleration scheme, which used two picosecond laser

pulses with a frequency difference equal to the plasma frequency. The interference of these long pulses resonantly excited a plasma wave capable of accelerating electrons [5, 6, 7].

With the rapid advancement of laser technology, and in particular the introduction of chirped pulse amplification (CPA) in 1985 by Strickland and Mourou [8], it became possible to generate ultrashort, high-peak-power femtosecond pulses. This technological breakthrough allowed the theoretical concept of laser wakefield acceleration (LWFA) to become experimentally feasible. Unlike beatwave schemes, a single sufficiently intense ultrashort pulse could directly drive a relativistic plasma wave without the need for two frequency-matched beams.

Early experiments in the late 1990s demonstrated multi-MeV electron acceleration using such CPA-based systems. However, the accelerated beams exhibited large energy spread and limited stability [9, 10, 11]. A major turning point occurred in 2004, when Faure et al. [12], and other groups [13, 14] reported the first quasi-monoenergetic electron beams, marking the emergence of LWFA as a viable compact accelerator concept.

In the following decade, many groups around the world achieved MeV–GeV electron beams with increasingly improved beam quality using LWFA-based concepts[15, 16, 17]. The current energy record was set in 2019 by Gonsalves et al., who reported 7.8 GeV electrons accelerated over a 20 cm long capillary[18].

Since the mid-2010s, the focus of the field has progressively shifted from demonstrating electron beam production to developing a viable accelerator technology. Despite the strong prospects of compact acceleration sources, many challenges remain before this concept can be used in practical applications. Conventional RF accelerators still offer superior reliability in terms of charge, beam quality, stability, and repetition rate.

To address these challenges, several research directions are currently being pursued. One major approach aims at improving the quality and control of the driving laser pulse. Another focuses on integrating advanced computational and machine learning methods to better control and optimize the electron beam properties. Recent advances include the RAVEN diagnostic in 2025, which allows single-shot full spatiotemporal vector field measurements of ultraintense laser pulses [19]. This promising diagnostic could enable feedback-loop laser control schemes.

At DESY, the LUX beamline has served as a testbed for advanced control methods in laser–plasma acceleration. In experiments reported by Jalas et al., Bayesian optimization was used to autonomously adjust laser and plasma parameters with the aim of reducing the longitudinal energy spread of electron bunches to sub-percent levels [20]. In a separate set of studies, extended operation of the LUX laser-plasma accelerator was reported. The LUX beamline was operated continuously for over a day, producing electron beams at approximately 1Hz for more than 30 hours, demonstrating sustained accelerator performance over extended periods [21]. More recently, DESY's

high-average-power laser KALDERA was commissioned. Initial experiments driving the MAGMA plasma accelerator achieved acceleration of up to 100 electron bunches per second [22], representing an increase in repetition rate compared with earlier laser-plasma accelerator prototypes. In parallel, another team demonstrated the first free-electron laser driven by a laser-plasma accelerator at Dresden, after meticulous improvements of the DRACO laser quality [23].

Finally, the PALLAS project is a laser–plasma injector hosted at the University of Paris-Saclay in Orsay and is part of the European EUPRAXIA initiative. Its goal is to advance plasma-based accelerator concepts and move toward making these accelerators a practical reality. Over the past three years, my work within PALLAS has focused on developing machine learning methods to improve electron beam quality and on implementing advanced laser diagnostics and algorithms to achieve better control over the driving laser pulse. I have also participated in experimental campaigns that led to the production of the first electron beams from the PALLAS platform.

This thesis is structured as follows: Chapter 1 establishes the theoretical foundations of laser wakefield acceleration (LWFA) and provides the necessary framework for subsequent discussions. Chapter 2 assesses the current operational status of the PALLAS test facility, with a focus on the first experimentally obtained electron beams. Chapter 3 explores the development and implementation of machine learning methodologies to enhance the performance and stability of the laser-plasma injector. In Chapter 4, advanced techniques for spatio-spectral laser pulse characterisation and phase retrieval are examined, addressing key challenges in the precise measurement of laser drivers. Chapter 5 presents a computational framework designed to model and predict the thermal deformation of laser compressors, a challenge for future high-repetition-rate laser-driven plasma accelerators. Finally, Chapter 6 concludes with a study of realistic laser-plasma acceleration simulations.

References

[1] John Douglas Cockcroft and Ernest T.S. Walton. "Experiments with high velocity positive ions.—(I) Further developments in the method of obtaining high velocity positive ions". In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 136.830 (1932), pp. 619–630.

[2] Helmut Wiedemann. *Particle Accelerator Physics*. 4th. Cham, Switzerland: Springer, 2015.

[3] ATLAS Collaboration. "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". In: *Physics Letters B* 716.1 (2012), pp. 1–29.

[4] T. Tajima and J.M. Dawson. "Laser electron accelerator". In: *Physical review letters* 43.4 (1979), p. 267.

[5] M. Everett et al. "Trapped electron acceleration by a laser-driven relativistic plasma wave". In: *Nature* 368.6471 (1994), pp. 527–529.

[6] A. E. Dangor, A. K. L. Dymoke-Bradshaw, and A. E. Dyson. "Observation of relativistic plasma waves generated by the beat-wave with 1 micron lasers". In: *Physica Scripta* T30 (1990), pp. 107–109.

[7] Y. Kitagawa et al. "Beat-wave excitation of plasma wave and observation of accelerated electrons". In: *Physical Review Letters* 68.1 (Jan. 1992), pp. 48–51.

[8] Donna Strickland and Gerard Mourou. "Compression of amplified chirped optical pulses". In: *Optics Communications* 56.3 (1985), pp. 219–221.

[9] C. A. Coverdale et al. "Propagation of intense subpicosecond laser pulses through underdense plasmas". In: *Physical Review Letters* 74.23 (June 1995), pp. 4659–4662.

[10] K. Nakajima et al. "Observation of ultrahigh gradient electron acceleration by a self-modulated intense short laser pulse". In: *Physical Review Letters* 74.23 (May 1995), pp. 4428–4431.

[11] A. Modena et al. "Electron acceleration from the breaking of relativistic plasma waves". In: *Nature* 377.6550 (1995), pp. 606–608.

[12] J. Faure et al. "A laser–plasma accelerator producing monoenergetic electron beams". In: *Nature* 431.7008 (2004), pp. 541–544.

[13] S. P. D. Mangles et al. “Monoenergetic beams of relativistic electrons from intense laser–plasma interactions”. In: *Nature* 431.7008 (2004), pp. 535–538.

[14] C. G. R. Geddes et al. “High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding”. In: *Nature* 431.7008 (2004), pp. 538–541.

[15] S. Kneip et al. “Near-GeV acceleration of electrons by a nonlinear plasma wave driven by a self-guided laser pulse”. In: *Physical Review Letters* 103.3 (July 2009), p. 035002.

[16] W. P. Leemans et al. “Multi-GeV electron beams from capillary-discharge-guided subpetawatt laser pulses in the self-trapping regime”. In: *Physical Review Letters* 113.24 (Dec. 2014), p. 245002.

[17] Nasr A. M. Hafz et al. “Stable generation of GeV-class electron beams from self-guided laser–plasma channels”. In: *Nature Photonics* 2 (Aug. 2008), pp. 571–577.

[18] A.J. Gonsalves et al. “Petawatt laser guiding and electron beam acceleration to 8 GeV in a laser-heated capillary discharge waveguide”. In: *Physical review letters* 122.8 (2019), p. 084801.

[19] S. Howard et al. “Single-shot spatiotemporal vector field measurements of petawatt laser pulses”. In: *Nature Photonics* 19.8 (2025), pp. 898–905.

[20] S. Jalas et al. “Bayesian optimization of a laser-plasma accelerator”. In: *Physical review letters* 126.10 (2021), p. 104801.

[21] Deutsches Elektronen-Synchrotron (DESY). *LUX Beamline: First Day-Long Operation of a Laser-Plasma Accelerator*. Press release. The LUX laser-plasma accelerator ran continuously for more than 24 hours, generating electron beams over extended periods. 2020. url: https://mpl.desy.de/lux.

[22] Deutsches Elektronen-Synchrotron (DESY). *DESY’s KALDERA Laser Achieves 100 Electron Bunches per Second*. Press release. The KALDERA laser drove the MAGMA plasma accelerator at up to 100 shots per second, opening the path to high-rep-rate plasma acceleration. 2025. url: https://desy.de/desy_latest_news/2025/100_shots_per_second/index_eng.html.

[23] M. Labat et al. “Seeded free-electron laser driven by a compact laser plasma accelerator”. In: *Nature Photonics* 17.2 (2023), pp. 150–156. doi: 10.1038/s41566-022-01104-w.

1 -Laser wakefield theory

1.1 Chirped pulse amplification

The Chirped pulse amplification (CPA) scheme, introduced by Gérard Mourou and Donna Strickland in 1985 [1], revolutionized laser technology by enabling systems to achieve peak powers in the Terawatt to Petawatt range. The scheme operates as follows: an oscillator with a broadband gain medium initially generates a low-energy femtosecond pulse in the nanojoule range. This pulse is then stretched to the picosecond range to prevent nonlinear effects in air and to avoid damage to optical components during amplification. After amplification, the pulse is subsequently compressed to recover its original femtosecond duration, producing an ultrafast, high-power laser pulse (Fig 1.1) .

In the following, we will briefly present each component of this laser chain, namely the oscillator, the stretcher, the amplification system, and the compressor.

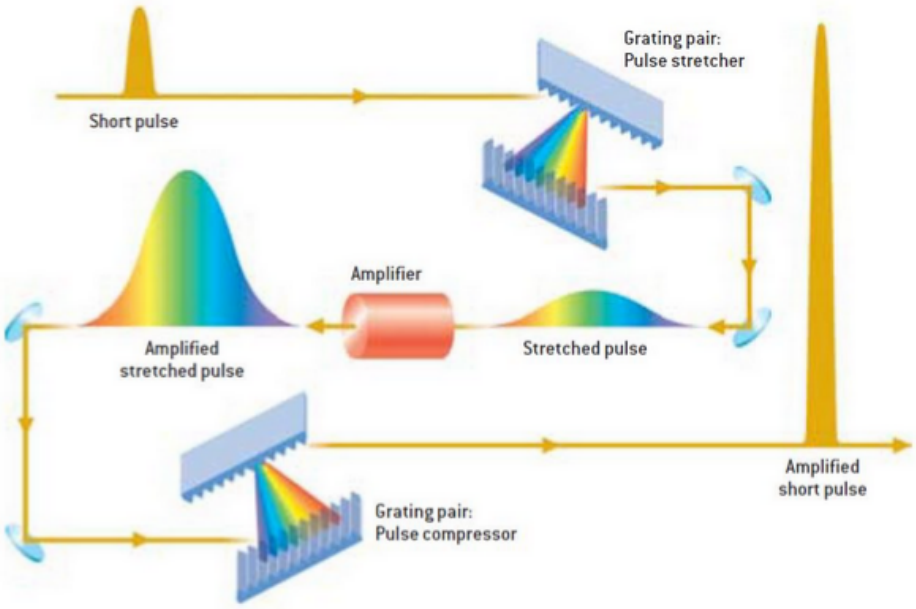

Figure 1.1: Simplified layout of the CPA laser system.

1.1.1 Oscillator

The laser oscillator can be modeled by a Fabry-Perot cavity inside of which resides a broadband gain medium which allows a large number of standing waves with slightly

different frequency to exist. To model the electric field inside the cavity, we can use the plane wave solution of the wave equation:

$$E(z,t) = e^{i(\omega t - kz + \phi)} \tag{1.1}$$

With ω the pulsation, k the wave number and ϕ a random phase term. To form a standing wave inside the cavity the phase must be the same after one round trip. If the cavity is of length L and with the light traveling at speed c we have:

$$\omega_n = \frac{n\pi c}{L}, \quad n \in \mathbb{N} \tag{1.2}$$

The frequency that can oscillate inside the cavity must be an integer multiple of the free spectral range (FSR), given by: $FSR = \frac{c}{2L}$. where c is the speed of light and L is the cavity length. Based on this, we can model the total electric field inside the cavity as the sum of all the standing waves:

$$E(t) \propto \sum_{n=-(N-1)/2}^{n=(N-1)/2} e^{i[(\omega_0 + n\Delta\omega)t + \phi_n]} \tag{1.3}$$

With $\Delta\omega = 2\pi FSR$, ω_0 the central pulsation and N the number of modes inside the cavity. While we cannot directly solve this sum because ϕ_n is a random variable that can take any value inside $[0, 2\pi]$ we can compute its expected value.

$$\mathbb{E}[E(t)] = \sum_{-(N-1)/2}^{(N-1)/2} \mathbb{E}[e^{i[(\omega_0 + n\Delta\omega)t + \phi_n]}] \tag{1.4}$$

This can be rewritten as:

$$\mathbb{E}[E(t)] = e^{i\omega_0 t} \sum_{-(N-1)/2}^{(N-1)/2} e^{in\Delta\omega t} \mathbb{E}[e^{i\phi_n}]] \tag{1.5}$$

If we assume ϕ_n is uniformly distributed within $[0, 2\pi]$, we have:

$$\mathbb{E}[e^{i\phi_n}] = \int_0^{2\pi} e^{i\phi} \frac{d\phi}{2\pi} = 0 \tag{1.6}$$

Thus, we conclude that the expected value of the total electric field is zero. We then compute the expected value of the square modulus of the electric field.

$$\mathbb{E}[|E(t)|^2] = \mathbb{E}\left[\sum_{n=-(N-1)/2}^{(N-1)/2} e^{i[(\omega_0 + n\Delta\omega)t + \phi_n]} \sum_{m=-(N-1)/2}^{(N-1)/2} e^{-i[(\omega_0 + m\Delta\omega)t + \phi_m]} \right] \tag{1.7}$$

This expression simplifies to:

$$\mathbb{E}[|E(t)|^2] = \sum_{n,m} e^{i[(n-m)\Delta\omega t]} \mathbb{E}\left[e^{i[\phi_n - \phi_m]}\right] \tag{1.8}$$

Now we can see from (1.8) that when $n = m$ the summand is equal to 1 and when $n \neq m$ the summand is equal to zero. Thus the expected value of the intensity is N. This implies that the standard deviation of the electric field is $\sqrt{N}$. Thus, at any given time t, we expect the modulus of the electric field to be proportional to the square root of the number of modes, with a variance also equal to the square root of the number of modes. This poses a problem if we want our oscillator to deliver femtosecond pulses. From this simple analysis, it becomes clear that in order to achieve such a pulse, it is necessary to introduce a constant phase relationship between the different modes of the oscillator. For simplicity we will assume $\phi_n = 0$ for all n. With this assumption we are then able to compute the intensity:

$$E(t) \propto \sum_{-(N-1)/2}^{(N-1)/2} e^{i(\omega_0 + n\Delta\omega)t} \tag{1.9}$$

We get the total intensity inside the cavity to be:

$$I(t) \propto \frac{\sin^2(\frac{N\Delta\omega t}{2})}{\sin^2(\frac{\Delta\omega t}{2})} \tag{1.10}$$

From equation (1.10), we can see that the resulting intensity forms an impulse comb with a period given by:

$$T = \frac{2\pi}{\Delta\omega} = \frac{1}{FSR}.$$

Additionally, the duration between the peak of a pulse and its first zero is given by:

$$\tau = \frac{2\pi}{N\Delta\omega}.$$

This shows that by introducing a fixed phase relationship between the different modes of the cavity, we can generate a pulse train with a period $T = \frac{1}{FSR}$ Furthermore, the pulse duration is inversely proportional to the number of modes inside the cavity. This explains why a broadband gain medium is essential for generating femtosecond pulses. A broader gain bandwidth supports a larger number of modes, leading to shorter pulse durations. In the following we will see how to introduce this fixed phase relationship between the different modes which is referred as mode-locking.

Mode locking

As we saw above, to generate a femtosecond pulse at the output of our oscillator, it is necessary to introduce a fixed phase relationship between the different modes. To

illustrate how this process is performed we will consider the temporal electric field of the mode n.

$$E_n(t) = E_0 \cos(\omega_n t + \phi_n)$$

Now assume that we are able to introduce inside the cavity a process which will periodically modulate the amplitude of this mode such that:

$$E_n = E_n(t)(1 + \alpha \cos(\Omega t + \phi))$$

Then the electric field of the mode n can be written as:

$$E_n = E_0 \left[\cos(\omega_n t + \phi_n) + \frac{\alpha}{2}\cos((\omega_n + \Omega)t + \phi_n + \phi) + \frac{\alpha}{2}\cos((\omega_n - \Omega)t + \phi_n - \phi)\right] \tag{1.11}$$

By introducing this amplitude modulation we have generated sidebands of frequency $\omega_n \pm \Omega$. Now if Ω is made such that it is equal to the mode frequency spacing , $\Delta\omega = 2\pi FSR$, The sidebands will be located at the adjacent modes $n \pm 1$. This coupling between adjacent modes results in the tendency of the different modes to lock their phases onto the sidebands. For example lets now consider that $\Omega = \Delta\omega$, if we now consider the mode n+1 and its side bands we get:

$$\alpha E_0 \left[\cos\left(\omega_{n+1}t + \frac{\phi_{n+2} + \phi_n}{2}\right)\cos\left(\frac{\phi_{n+2} - (\phi_n + 2\phi)}{2}\right)\right] + E_0 \cos(\omega_{n+1}t + \phi_{n+1}) \tag{1.12}$$

The modulus of this term is maximised for $\phi_{n+2} - \phi_n = 4p\pi + 2\phi$ and $\phi_{n+1} - \phi_n = 2p\pi + \phi$,$p \in \mathbb{N}$. With this simple example we can understand that if the gain of the cavity is modulated at the frequency of the free spectral range then modes that exhibit a constant phase relationship with their neighboring modes will be favored in the amplification competition. If this condition is applied for several round trips in the cavity eventually all modes will be locked in phase and a single pulse containing all the energy will appear. This approach is referred to as amplitude modulation mode locking, this is not the only approach to introduce a coupling between the different modes. There is another method called frequency modulation mode locking[2, p. 388] that we will not see here where the idea is to periodically modulate the frequency of the electric field in order to introduce a coupling between modes.

To be analysed in detail the mode locking mechanisms need a more thorough investigation. Indeed the simple picture that we gave above is far from complete. In fact one of the main problem is that we assumed that the mode spacing of the cavity was constant across all frequencies. In reality the gain medium in the cavity is inherently dispersive thus the mode spacing is not constant. Moreover the laser cavity is not perfectly rigid thus it is hard to talk about a fixed number of modes. In reality this is during the

mode locking mechanism itself that the frequency comb will form through an interplay between self phase modulation in the gain medium and dispersion mechanisms. In the following I will present succinctly one of the more commonly used method of mode locking, the Kerr lens mode locking (KLM).

Kerr lens mode locking (KLM)

The KLM relies on a third order non linear effect in $\chi^{(3)}$ materials where the index of refraction of the material is dependent on the intensity of the light that passes through it.

$$n(r,t) = n_0 + n_2 I(r,t) \tag{1.13}$$

This property of the material leads to two discernable effects: the first one is self phase modulation the second is self-lensing.

Self phase modulation

Consider a light beam progressing in the (OZ) direction with both temporal and spatial gaussian profiles:

$$I(r,t) = I_0 e^{-2\frac{t^2}{\tau^2}} e^{-2\frac{r^2}{w(z)^2}}$$

The phase inprit of such a beam in a $\chi^{(3)}$ material of length L would be:

$$\phi(r,t) = \omega_0 t - kL(n_0 + n_2 I(r,t)) \tag{1.14}$$

With k the wave number $k = \frac{\omega_0}{c}$. Now if we compute the derivative of this phase with respect to time we obtain the instantaneous frequency inside the pulse:

$$\omega(r,t) = \omega_0 + \frac{4kLn_2 I(r,t)t}{\tau^2}$$

The second term in the right hand side can be approximated by a linear term that is reffered to as the chirp of the beam. This term represents the shift from the carrier frequency that occurs inside the beam. Since n_2 is a positive number the leading part of the pulse experiences lower frequencies than the tail. The new frequencies generated in the light beam, broaden the frequency spectrum of the beam. This effect is important in KLM because if the spectral phase distortions introduced by dispersion and self-phase modulation are properly compensated, the shortest possible pulse duration can be achieved. In FIG. 1.2 we give an illustration of the SPM mechanism

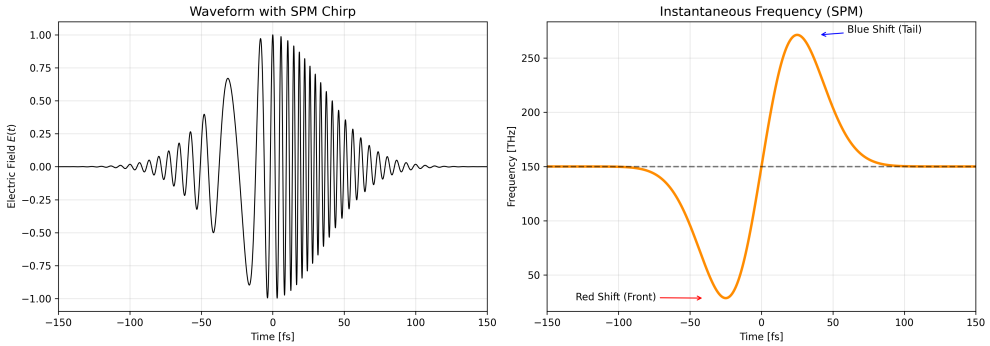

Figure 1.2: Illustration of the SPM mechanism for $\lambda = 2\,\mu\mathrm{m}$ (for visualization purposes), $n_2 = 2 \times 10^{-20}\,\mathrm{m}^2/\mathrm{W}$, $L = 1\,\mathrm{cm}$, and $\tau = 50\,\mathrm{fs}$. Left panel: chirped electric field. Right panel: instantaneous frequency.

Self lensing

To introduce the self lensing effect we will start with the phase (1.14). At the center part of the beam, we can approximate the tranverse intensity by a parabolic curve. The phase then becomes:

$$\phi(r,t) \simeq \omega_0 t - kL\left(n_0 + n_2 I(t)\left(1 - \frac{2r^2}{w_0^2}\right)\right) \tag{1.15}$$

If we focus on the quadratic term of the phase we can see that it looks similar to the phase factor of a thin lens. Indeed for a thin lens we have:

$$T(r) = \exp\left(\frac{-ikr^2}{2f}\right)$$

Thus we can see that the Kerr medium acts as a thin lens with a time dependent focal length:

$$f(t) = \frac{w_0^2}{4Ln_2 I(t)}$$

When the light beam is propagating through the Kerr medium it undergoes a dynamical lens whose vergence is increasing with the light intensity. This effect is at the heart of KLM because this is the one who provides the amplitude modulation needed to sustain the mode locking.

Once the laser is pumped and begins operating, it initially starts in a continuous wave regime. However, if a sufficiently intense pulse forms due to fluctuations, the Kerr lensing effect induces an intensity-dependent focusing of the beam. When an appropriately placed aperture is introduced in the cavity, this focusing results in an effective amplitude

modulation: higher-intensity pulses are more tightly focused and pass through the aperture with less losses, while lower-intensity fluctuations experience greater diffraction and attenuation. This preferential amplification of the most intense pulses reinforces phase locking between the different modes, as phase-coherent modes constructively interfere to form the strongest pulse.

Since this amplitude modulation occurs at the round-trip time of the cavity, it satisfies the condition for phase locking discussed earlier. As a result, over multiple round trips, the different modes progressively synchronise, and a single dominant pulse emerges, capturing nearly all the available energy.

Simultaneously, self-phase modulation (SPM) broadens the pulse spectrum due to the intensity-dependent refractive index of the gain medium. However, dispersion from the optical elements in the cavity causes temporal broadening, which must be compensated to maintain ultrashort pulse durations. This is achieved using a dispersion-compensating system, typically a set of prisms or diffraction gratings, which balances the group delay dispersion introduced by SPM and the cavity itself. Once this balance is reached, the oscillator outputs a stable, ultrashort femtosecond pulse. In FIG. 1.3 we give an illustration of the KLM mechanism

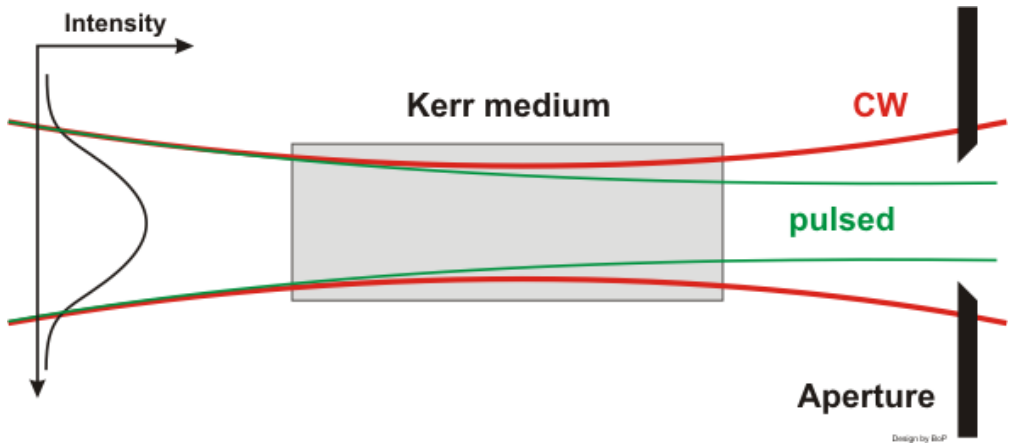

Figure 1.3: Illustration of the KLM mechanism. The continuous-wave modes are blocked by the aperture, while the pulsed mode is continuously reinforced at each round trip inside the cavity. Figure adapted from Wikipedia.

1.1.2 Stretcher and compressor

Once the fs pulse exits the oscillator it is sent to the stretcher in order to be temporally expanded and thus prevent any optical damages in the laser chain after amplification. To understand how the stretcher operates, we first need to understand how the spectral

phase impacts the duration of the beam. Consider a spectrally Gaussian electric field with a non zero second order phase term (chirp):

$$E(\omega) = E_0 e^{(-a(\omega-\omega_0)^2 + ib(\omega-\omega_0)^2)} \tag{1.16}$$

By computing the Fourier transform of this field we find the temporal electric field as:

$$E(t) = \mathscr{F}\left\{E(\omega)\right\} = E_0\sqrt{\frac{\pi}{a-ib}}\exp\left(\frac{-t^2}{4(a-ib)} - i\omega_0 t\right) \tag{1.17}$$

We thus have for the intensity:

$$I(t) = \frac{\pi E_0^2}{\sqrt{(a^2+b^2)}}\exp\left(\frac{-at^2}{2(a^2+b^2)}\right) \tag{1.18}$$

this expression leads us to:

$$\frac{a}{2(a^2+b^2)} = \frac{1}{\tau^2}$$

$$\tau = \sqrt{2a}\sqrt{1+\left(\frac{b}{a}\right)^2}$$

With τ the pulse duration and $a = \frac{1}{2\Delta\omega^2}$, $\Delta\omega$ the spectral bandwidth. With this expression we see that if we consider a pulse with spectral chirp the temporal duration will increase by a factor $\sqrt{1+\left(\frac{b}{a}\right)^2}$.

More generally, since the spectral phase of lasers is a narrowband signal, it can usually be expressed as a Taylor serie expansion:

$$\Phi(\omega) = \phi^{(0)} + \phi^{(1)}(\omega-\omega_0) + \frac{\phi^{(2)}}{2}(\omega-\omega_0)^2 + ... + \frac{\phi^{(n)}}{n!}(\omega-\omega_0)^n \tag{1.19}$$

with $\phi^{(0)} = \phi(\omega=\omega_0)$ and $\phi^{(i)} = \left(\frac{\mathrm{d}^i\Phi}{\mathrm{d}\omega^i}\right)_{\omega=\omega_0}$

This is convenient because it allows us to isolate and describe the impact of each term in the series as a distinct physical effect, making it easier to understand and manage the contributions to the overall spectral phase in practical systems. We will say a few words on the first four terms of the expansion.

- $\phi^{(0)}$ is the carrier envelope phase, it describes the relative phase between the envelope and the carrier.
- $\phi^{(1)}$ only represents the temporal delay between the center of the pulse and an arbitrary time origin, it is refereed to as group delay.

- $\phi^{(2)}$ referred to as group delay dispersion (GDD) or chirp and represents the rate at which the frequencies within the pulse deviate from the central frequency. As we saw above, this term plays a crucial role, since it can cause temporal broadening of pulses.
- $\phi^{(3)}$ This term is referred as third order dispersion (TOD), it introduces asymmetry in the pulse by generating temporal distortions, leading to the formation of pre-pulses or post-pulses.

Since the signal is narrowband, higher-order dispersion terms have a lesser impact, provided that the pulse does not undergo strong nonlinear effects. In practice, this allows the Taylor expansion of the spectral phase to be truncated at the third-order term. The stretcher is designed to introduce a controlled amount of GDD, while the compressor compensates for both GDD and TOD, ensuring the pulse recompression to its shortest possible duration.

The typical compressor in a femtosecond laser system is composed of a pair of two parallel gratings spaced by a distance L. To understand how this system can introduce GDD in the pulse we will consider the grating equation:

$$\sin\alpha - \sin\beta(\omega) = \frac{2\pi nc}{\omega d}, \quad n \in \mathbb{Z} \tag{1.20}$$

where α is the incidence angle, β is the diffraction angle, d is the groove spacing of the grating, and n is the order of diffraction. This equation describes the relationship between the incidence and diffraction angles for a given frequency and grating period. In what follows, we will consider only the first diffraction order.

Referring to Fig. 1.4, we can determine the path length ABC and thus the accumulated phase of a light ray passing through a pair of gratings:

$$\Phi(\omega) = \frac{\omega L}{c\,\cos\beta(\omega)}(1 + \cos(\alpha + \beta(\omega))) \tag{1.21}$$

Using equation (1.20), the first derivative is given by:

$$\frac{\mathrm{d}\Phi}{\mathrm{d}\omega} = \left(\frac{L}{c}\right)\frac{1 + \cos(\alpha + \beta(\omega))}{\cos(\beta(\omega))} + \frac{2\pi L}{d\,\cos^2(\beta(\omega))}\frac{\mathrm{d}\beta}{\mathrm{d}\omega} \tag{1.22}$$

For the second term on the right-hand side, we find its primitive:

$$\frac{2\pi L}{d}\tan(\beta(\omega)) \tag{1.23}$$

From Fig. 1.4, we see that this term corresponds to a 2π multiple of the number of grating grooves between point B and the intersection of the normal to G_1 on G_2. This reflects the periodic nature of the grating.

If two light rays of the same frequency strike G_1 at different locations, they will exit G_1 parallel to each other and belong to the same wavefront. After diffraction at G_2,

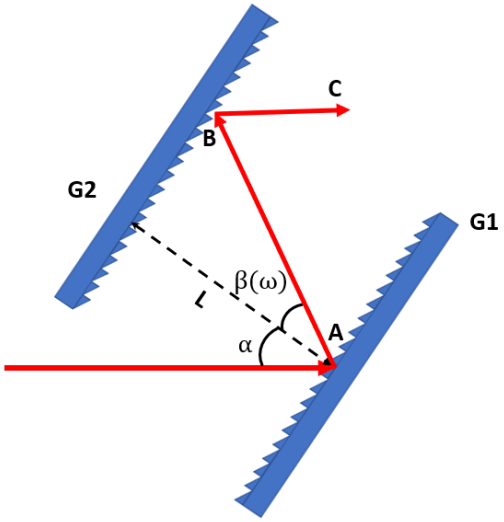

Figure 1.4: Sketch of a compressor.

these two rays will still be part of the same wavefront, even though their path lengths from G_1 to the output wavefront are different. This is because, although one ray travels a longer distance than the other, their optical path difference is always an integer multiple of the wavelength, corresponding to the number of grooves separating them on G_2. As a result, their relative phase remains the same, while their absolute phase is shifted by an integer multiple of 2π.

If $\beta(\omega)$ increases so that a light ray now strikes an adjacent groove, the accumulated phase $\Phi(\omega)$ undergoes a 2π phase jump to maintain the phase relation between parallel rays. Since we are only concerned with relative phase, this term can be canceled by adding its negative counterpart in equation (1.21), leading to the simplified first derivative:

$$\frac{\mathrm{d}\Phi}{\mathrm{d}\omega} = \left(\frac{L}{c}\right) \frac{1+\cos(\alpha+\beta(\omega))}{\cos(\beta(\omega))} \tag{1.24}$$

The second order derivative yields

$$\frac{\mathrm{d}^2\Phi}{\mathrm{d}\omega^2} = \frac{-4\pi^2 Lc}{\omega^3 d\cos^3(\beta(\omega))} \tag{1.25}$$

An interesting feature of equation (1.25) is that it is always negative. the third derivatives yields:

$$\frac{\mathrm{d}^3\Phi}{\mathrm{d}\omega^3} = \frac{-3}{\omega}\left[1 - \frac{2\pi c\sin(\beta(\omega))}{\omega d\cos^2(\beta(\omega))}\right]\frac{\mathrm{d}^2\Phi}{\mathrm{d}\omega^2} \tag{1.26}$$

So in order to tune the amount of GDD introduced by the compressor one can play with the distance L separating the gratings. In order to tune the ratio between GDD and TOD, one can play with the incidence angle and the groove spacing of the gratings.

The stretcher in typical laser systems presents a similar geometry but the gratings are separated by an imaging system in order to introduce some positive GDD. In the end the compressor is tailored such that it compensates most of the positive GDD introduced by the stretcher and the optics.

1.1.3 Amplification

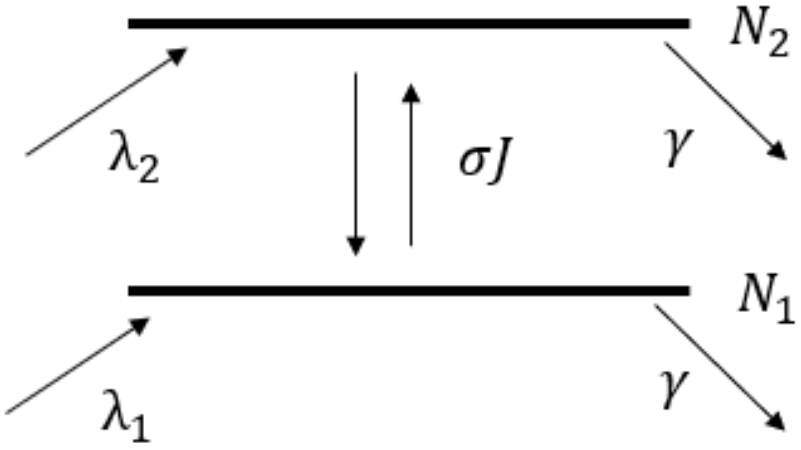

Figure 1.5: Sketch of a two-level system

The amplification process is at the heart of every laser system and is based on the principle of stimulated emission of radiation. The principle is as follows: a gain medium, located inside an optical cavity, is pumped by an external energy source that excites the atoms of the medium, lifting their population from the ground state to a higher energy level. When a photon with the proper energy $h\nu$ [J] passes through the gain medium, there is a probability that an excited atom will de-excite, emitting another photon of the same energy and phase, coherent with the first one. By exploiting this process, the field intensity can grow exponentially until it reaches saturation.

To illustrate this process, we consider a simple model of a two-level system. Although such a system is physically unrealistic, it allows us to identify the qualitative properties of the amplification phenomenon. The two-level system is represented schematically in Fig. 1.5. In this model we neglect spontaneous emission. The parameters are

defined as follows: λ_1 [s^{-1}] and λ_2 [s^{-1}] are the pumping rates into level 1 and level 2, respectively; γ [s^{-1}] is the non-radiative de-excitation rate; σ [m^2] is the stimulated emission cross section, representing the probability of interaction of incoming photons; J [$\mathrm{m}^{-2}\,\mathrm{s}^{-1}$] is the photon flux; finally, N_1 [m^{-3}] and N_2 [m^{-3}] denote the population densities of levels 1 and 2, respectively.

In this system, the evolution of the photon flux is described by:

$$\frac{dJ}{dt} = -kJ + c\sigma J\Delta N, \tag{1.27}$$

where k is the loss parameter, representing photons lost during the amplification process (primarily due to mirror transmission and scattering), and $\Delta N = N_2 - N_1$ is the population inversion.

From this equation, we see that amplification is possible only if the right-hand side is positive, i.e. only if $\Delta N > 0$. Furthermore, amplification requires:

$$\Delta N > \frac{k}{c\sigma} \equiv \Delta N_s,$$

where ΔN_s is the threshold population inversion, below which amplification is not possible. Since k represents the losses, the term $c\sigma\Delta N$ can be interpreted as the effective gain of the cavity.

The dynamics of the population inversion are given by the rate equation:

$$\frac{d\Delta N}{dt} = -\gamma(\Delta N - \Delta N_0) - 2\sigma J\Delta N, \tag{1.28}$$

where

$$\Delta N_0 = \frac{\lambda_2 - \lambda_1}{\gamma}$$

is the equilibrium inversion in the absence of incoming photons. The first term on the right-hand side drives the system toward its equilibrium value ΔN_0, while the second term accounts for stimulated emission when $\Delta N > 0$.

The steady-state solutions of the system are:

$$\begin{cases} J = 0, & \dfrac{\Delta N}{\Delta N_s} = A, \\ J = \dfrac{\gamma}{2\sigma}(A-1), & \dfrac{\Delta N}{\Delta N_s} = 1, \end{cases}$$

where $A = \dfrac{\Delta N_0}{\Delta N_s}$ is the pumping parameter.

The first solution corresponds to the case without photons in the cavity: the pumping simply drives the inversion to its equilibrium value ΔN_0. If $\Delta N_0 < \Delta N_s$, amplification is not possible because losses dominate. The second solution corresponds to the

case where the gain exactly compensates the losses: here, the photon flux is proportional to the pumping rate, and the population inversion is equal to the threshold value ΔN_s.

The pumping parameter can also be expressed in terms of the photon flux relative to the saturation flux:

$$A = 1 + \frac{J}{J_s}, \qquad J_s = \frac{\gamma}{2\sigma},$$

where J_s is the saturation flux,which represents the photon flux that makes the gain divided by two. In this form, the gain can be written as

$$c\sigma\Delta N = \frac{c\sigma\Delta N_0}{1 + \frac{J}{J_s}}.$$

This expression shows that as the photon flux in the cavity increases, the gain decreases (gain saturation). For large incoming photon flux, the gain approaches zero, and only marginal amplification remains possible.

In practice, TW and PW laser systems employ multiple amplification stages to bring the pulse energy from the nJ level of the seed to the multi-joule regime. Each stage is designed with an appropriate crystal size and pump energy so that the pulse fluence reaches the saturation fluence, ensuring efficient energy extraction. By stepping through successive stages in this way, the system maintains spectral quality while scaling the pulse energy to the level required for compression to TW or PW peak powers (see Fig. 1.6).

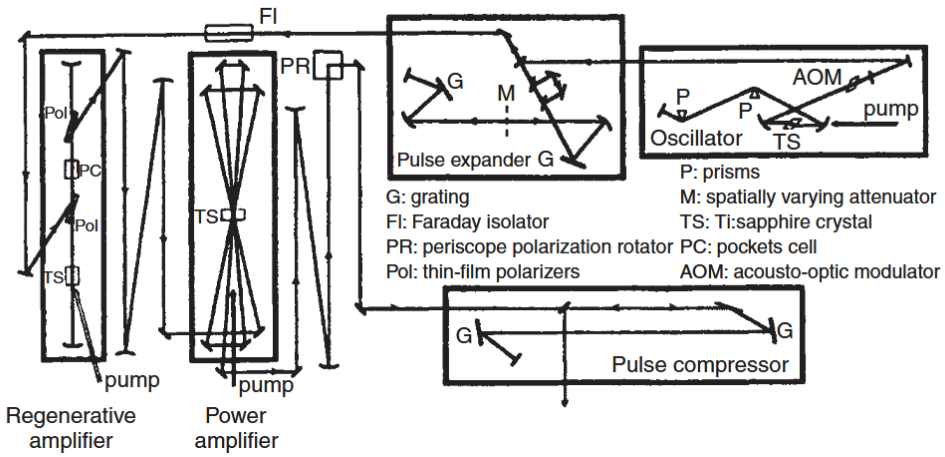

Figure 1.6: Sketch of a Multiterawatt fs laser system. (Adapted from [3])

1.2 Gaussian beam formalism

An electromagnetic wave propagating in vacuum is described by Maxwell's equations from which we can easily establish a wave equation for the electric field:

$$\Delta E - \frac{1}{c^2}\frac{\partial^2 E}{\partial t^2} = 0 \tag{1.29}$$

If we consider an harmonic solution of the equation (1.29), i.e., $E = u(x,y,z)e^{i(\omega t - kz)}$ and if we are interested only in slowly varying enveloppe solutions,ie:

$$\left|\frac{\partial u}{\partial z}\right| \ll |ku| \quad and \quad \left|\frac{\partial^2 u}{\partial z^2}\right| \ll \left|k\frac{\partial u}{\partial z}\right|$$

We can arrive at the paraxial wave equation:

$$\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} - 2ik\frac{\partial u}{\partial z} = 0 \tag{1.30}$$

A solution is the Gaussian beam:

$$E(x,y,z,t) = \sqrt{\frac{2}{\pi}}\frac{E_0 w_0}{w(z)}\exp\left\{\left[\frac{-(x^2+y^2)}{w(z)^2} - ik\frac{(x^2+y^2)}{2R(z)}\right]\right\}e^{i(\omega_0 t - k_0 z + \phi(z))} \tag{1.31}$$

This solution describes a beam propagating in the (OZ) direction with a transverse size that increases with propagation. It also captures the parabolic shape of the wavefront. The parameters of this equation are the following:

$$\begin{cases} z_R = \frac{\pi w_0^2}{\lambda} \\ w(z) = w_0\sqrt{1+\left(\frac{z}{z_R}\right)^2} \\ R(z) = z + \frac{z_R^2}{z} \\ \phi(z) = \arctan\left(\frac{z}{z_R}\right) \end{cases}$$

z_R is the Rayleigh length, it is a characteristic propagation for which the transverse area of the beam is doubled from the minimum size. $w(z)$ is the radius of the beam defined as $1/e^2$ of the maximum, its minimum value w_0 is called the waist and is at the position of the focus. $\phi(z)$ is the Gouy phase. This term represents the phase shift that the Gaussian wave undergoes in the vicinity of the focus. This term varies between $-\pi/2$ and $\pi/2$ and its value is almost constant far from the Rayleigh range. A more general solution to the paraxial wave equation is the Hermite-Gaussian (HG) beam, which is expressed as:

$$\begin{aligned} E_{m,n}(x,y,z,t) = \sqrt{\frac{2}{\pi}}\frac{E_0 w_0}{w(z)} H_m\left(\sqrt{2}\frac{y}{w(z)}\right) H_n\left(\sqrt{2}\frac{x}{w(z)}\right) \\ \times \exp\left[\frac{-(x^2+y^2)}{w(z)^2} - ik\frac{(x^2+y^2)}{2R(z)}\right] e^{i(\omega t - kz + \phi(z))}. \end{aligned} \tag{1.32}$$

These transverse modes typically arise in laser cavities with rectangular symmetry and are especially noticeable when the system is misaligned. In a well-aligned laser cavity, however, most of the energy is concentrated in the fundamental Gaussian mode TEM_{00}, which has the lowest losses and best beam quality. Higher-order HG modes appear when the cavity mirrors or gain medium introduce astigmatism or asymmetry, causing the beam to deviate from a pure Gaussian shape. The functions H_m and H_n represent Hermite polynomials, which describe the transverse intensity distribution along the x and y directions.

While the beam waist size $w(z)$ and radius of curvature $R(z)$ remain unchanged in the Hermite-Gaussian framework, the Gouy phase shift is modified as:

$$\phi(z) = (1 + m + n) \arctan\left(\frac{z}{z_R}\right).$$

Another important family of solutions is the Laguerre-Gaussian (LG) modes, which naturally arise in systems with cylindrical symmetry, such as ring resonators and fiber lasers. Like HG modes, LG modes are typically observed in misaligned cavities or when intentional mode conversion is applied.

In equation (1.31), we assumed that the temporal component of the laser beam is a pure sine wave, which is unrealistic since it would require a gain medium with infinite bandwidth. As mentioned earlier, the spectrum of laser pulses in fs systems is typically Gaussian in shape, meaning that the temporal profile is also Gaussian. This leads to the Fourier-limited Gaussian pulse:

$$E(x,y,z,t) = \sqrt{\frac{2}{\pi}} \frac{E_0 w_0}{w(z)} \exp\left\{\left[-2\ln 2 \frac{(z-ct)^2}{c\tau_0^2}\right]\right\} \exp\left\{\left[\frac{-(x^2+y^2)}{w(z)^2} - ik\frac{(x^2+y^2)}{2R(z)}\right]\right\} e^{i(\omega_0 t - k_0 z + \phi(z))} \tag{1.33}$$

with τ_0 the full width half maximum (FWHM) pulse duration. In practical laser systems the laser pulse often displays a super gaussian shape instead of a simple gaussian where the super gaussian would be defined as:

$$E(x,y,z,t) = \sqrt{\frac{2}{\pi}} \frac{E_0 w_0}{w(z)} \exp\left[-\alpha\left(\frac{z-ct}{c\tau_0}\right)^{2n}\right] \exp\left[-\left(\frac{x^2+y^2}{w(z)^2}\right)^m - ik\frac{x^2+y^2}{2R(z)}\right] e^{i(\omega_0 t - k_0 z + \phi(z))} \tag{1.34}$$

With α being a constant that ensures τ_0 represents the FWHM duration, and n and m being the orders of the temporal and spatial super-Gaussian profiles, respectively. The super-Gaussian shape arises in fs laser systems due to several phenomena. One of which is the saturation of the gain medium. During the amplification stage in the CPA chain, the central part of the pulse, which contains most of the energy, saturates the gain medium more quickly than the outer tails. As a result, it experiences less amplification, and after multiple round trips, this effect tends to flatten the shape of the pulse.

Another reason for the super-Gaussian shape is the homogenisation of the pump laser. To extract most of the energy from the gain medium, the shape of the pump is flattened using homogenisers, which triggers population inversion at the edges of the amplification crystal.

1.3 Optical aberrations

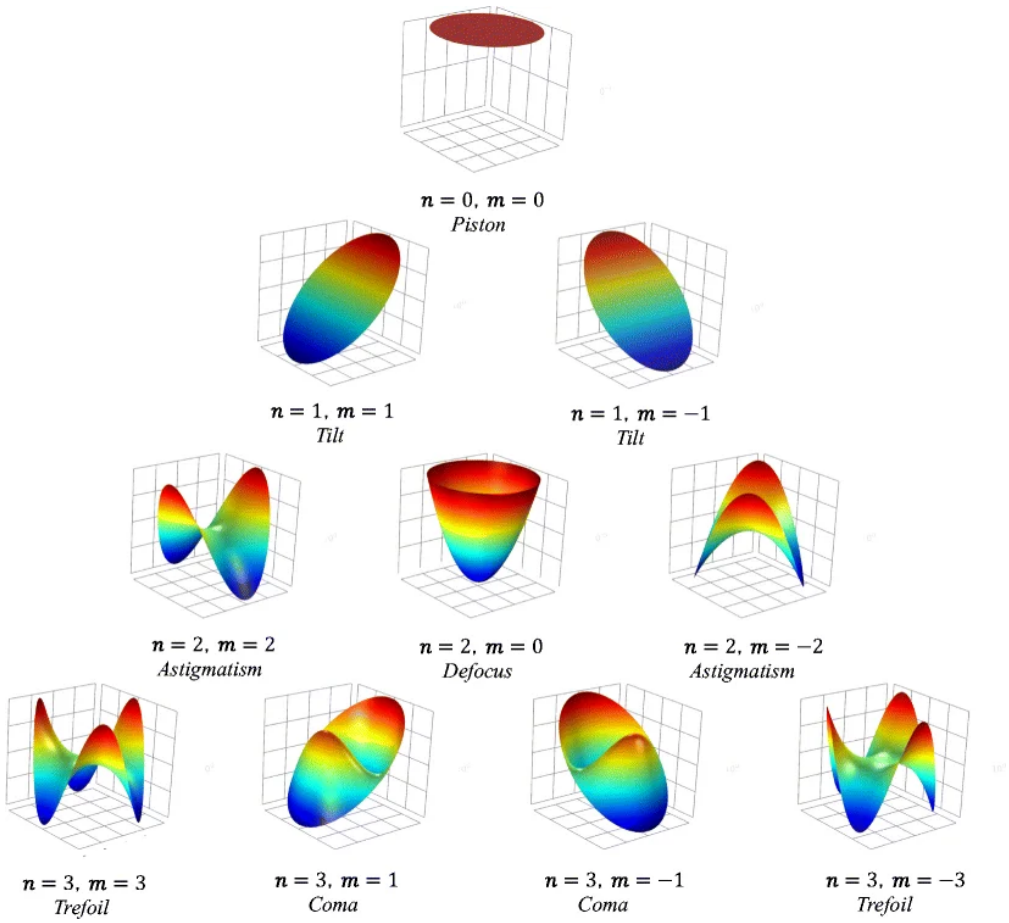

Figure 1.7: Visualization of the first Zernike polynomial modes representing common wavefront aberrations. Adapted from [4].

The laser beam used in the laser plasma interaction is subject to optical aberrations which can have an impact on the electron beam properties by modifying the intensity distribution at focus. Aberrations can be defined as the optical path difference from an ideal wavefront, where the ideal wavefront is the phase term of the Gaussian laser beam. These aberrations arise from imperfections in optical components,

thermal/mechanical deformations or misalignment. I will briefly present the most commonly found aberration in a CPA Laser chain.

- Spherical Aberration: This aberration arises in optical components with spherical surfaces, where rays incident at different transverse positions do not converge to the same focal point. As a result, the focal point becomes elongated, reducing the peak intensity and degrading the imaging quality. In high-intensity laser applications, spherical aberration can significantly affect the ability to achieve a tight, diffraction-limited focus.
- Astigmatism: Astigmatism occurs when an optical component lacks cylindrical symmetry, either due to manufacturing imperfections or misalignment. In such cases, the beam encounters different radii of curvature along two orthogonal planes: the sagittal plane and the tangential plane. Consequently, the beam focuses at two distinct focal planes, leading to an elongated or elliptical focal spot instead of a single sharp focus. This aberration is particularly detrimental in laser-plasma acceleration, where precise beam quality is crucial.
- Coma: Coma arises either due to the inherent design of certain optical elements or as a result of misalignment. It is an aberration that primarily affects off-axis beams, leading to an asymmetrical, comet-like distortion of the focal spot. It occurs when different regions of the optical element introduce varying magnifications across the aperture, causing the beam to stretch asymmetrically in one direction.
- Defocus: Defocus is a fundamental aberration that occurs when the longitudinal focal position does not coincide with the actual focus of the optical system. Instead of forming a sharp, well-defined spot, the beam appears blurred or larger than intended. Defocus is one of the simplest aberrations to correct since it just requires to modify the optical path length.

The aberrations discussed above are the primary distortions commonly encountered in optical systems. In a typical laser setup without spatial phase control mechanisms, the beam often exhibits a combination of these basic aberrations, along with more complex, higher-order distortions that further degrade the beam quality. A convenient mathematical framework used to describe such aberrations is the Zernike Polynomials. Zernike Polynomials are a set of orthogonal polynomials defined on a unit circle (see FIG. 1.7), with the following definition:

$$Z(r,\theta)_n^m = R(r)_n^m G(\theta) \tag{1.35}$$

$$\begin{cases} Z(r,\theta)_n^m = R(r)_n^m \cos(m\theta) & even \\ Z(r,\theta)_n^{-m} = R(r)_n^m \sin(m\theta) & odd \end{cases}$$

Here r and θ are the polar coordinates on the unit circle. With $n \geq m \geq 0$

$$R(r)_n^m = \begin{cases} \sum_{k=0}^{\frac{n-m}{2}} \frac{(-1)^k (n-k)!}{k!(\frac{n+m}{2}-k)!(\frac{n-m}{2}-k)!} r^{n-2k} & if \quad n-m \quad even \\ 0 \quad if \quad n-m \quad odd \end{cases} \tag{1.36}$$

Any sufficiently smooth wavefront can be represented as a linear combination of a limited number of Zernike polynomials where each (n,m) couple can be related to a specific aberration.

$$W(r,\theta) = \sum_{n,m} A_{m,n} Z(r,\theta)_n^m + B_{m,n} Z(r,\theta)_n^{-m} \tag{1.37}$$

With W the wavefront function, $A_{m,n}$ and $B_{m,n}$ the coefficients for the even and odd polynomials respectively. The truncation of the above sum is made such that the difference between the actual and reconstructed wavefront is minimised.

To assess beam quality and optical aberrations, several useful metrics exist. One of the simplest way to evaluate the wavefront is the peak-to-valley (PV) function, which measures the path length difference between the highest (most convex) and lowest (most concave) points on the wavefront. However, this metric can be misleading, as it only considers two sample points and is highly sensitive to noise. Additionally, it provides no information about the overall wavefront shape. A more robust quantity is the root mean square (RMS) error, which accounts for the entire wavefront surface and provides a statistically meaningful representation of wavefront deviations. The definition of RMS error:

$$\sigma = \sqrt{\langle \Delta W^2 \rangle - \langle \Delta W \rangle^2} \tag{1.38}$$

With $\langle . \rangle$ the mean ,ΔW the wavefront error (departure from the ideal Gaussian wavefront). Another important quantity in optical imaging is the *Strehl ratio* (SR), defined as the ratio of the peak intensity at the focal point between an aberrated wavefront and an unaberrated one. In the image plane, an unaberrated beam produces an Airy pattern with its maximum intensity at the center. Using the Fraunhofer diffraction integral, the Strehl ratio can be expressed as:

$$SR = \frac{1}{\pi^2} \left| \int_S e^{ik\Delta W} \, r \, dr \, d\theta \right|^2, \tag{1.39}$$

where ΔW is the wavefront error.

Expanding $e^{ik\Delta W}$ in a Taylor series and neglecting terms of third order and higher, the Strehl ratio can be approximated as

$$SR \simeq 1 - k^2 \sigma^2,$$

A more robust approximation, valid under broader conditions, is given by the Maréchal expression:

$$SR \simeq e^{-k^2\sigma^2}. \tag{1.40}$$

A SR of 1 corresponds to a perfect, unaberrated system. As a rule of thumb, the Maréchal criterion considers an optical system to be *well corrected* when $SR \geq 0.8$.

—

We now turn to the *point spread function* (PSF), which is the impulse response of the imaging system. The PSF is a two-dimensional function describing the image formed by the system from a point source. For a perfect optical system, the PSF corresponds to an Airy disk.

Mathematically, the PSF can be expressed as

$$PSF(\rho,\varphi) \propto \left| \int_S e^{ik\Delta W(r,\theta)} e^{i\frac{k}{f} r\rho\cos(\theta-\varphi)}\, r\, dr\, d\theta \right|^2, \tag{1.41}$$

where (r,θ) are pupil-plane polar coordinates and (ρ,φ) are image-plane coordinates.

Because imaging systems are linear and shift-invariant, the image of an arbitrary object O can be obtained by the convolution of the object with the PSF:

$$I(\rho,\varphi) = (O * PSF)(\rho,\varphi).$$

Here, $I(\rho,\varphi)$ is the recorded image, O is the object, and $*$ denotes the convolution operation.

1.4 Spatio-temporal/spectral couplings

In the previous section, we considered only spatial aberrations that can affect the laser beam. However, ultrashort laser pulses are broadband in frequency and are therefore also subject to chromatic aberrations. These aberrations arise from the frequency-dependent nature of light propagation. A simple example is the focusing of a broadband pulse through a lens: because the refractive index of the lens depends on frequency, each spectral component is focused at a different longitudinal position. As a result, the focal spot acquires a spatio-spectral dependence. Since the temporal and spectral properties of a laser pulse are connected through the Fourier transform, this spatio-spectral dependence generally leads to spatio-temporal couplings (STC). Depending on the goal of the experiment STC can be either desirable or detrimental. A more formal definition of STC is the following:

$$E(x,y,t) \neq f(x,y)g(t) \quad \forall \quad f(x,y), g(t) \tag{1.42}$$

It is an electric field that cannot be written as the product of a function of time and a function of space. One of the main STC is known as pulse front tilt (PFT),(see FIG. 1.8a).

This occurs when a laser pulse's main axis is tilted with respect to the direction of propagation, meaning that the arrival time of the laser varies linearly across the transverse direction. One of the main causes of PFT is angular dispersion (AD). AD is another form of spatio-spectral coupling (SSC), in which different frequency components propagate at different angles. A primary source of AD is the misalignment or deformation of the pulse compressor. AD is also the main origin of spatial chirp (SC),(see FIG. 1.8b), which corresponds to a transverse displacement of the different frequencies composing the pulse. In other words, the spectral components of the pulse vary linearly with the transverse position.

In the following, we present a general formalism for first-order STC [5] that provides a unified framework for understanding the relationships between the different types of couplings.

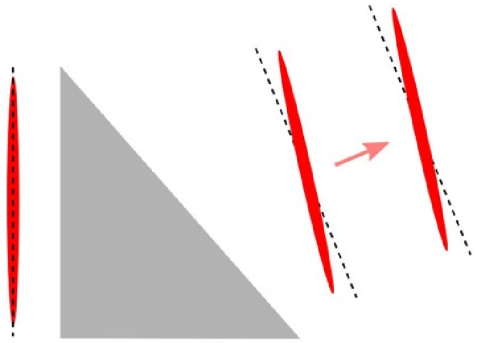

(a) Pulse Front Tilt (PFT) induced by angular dispersion (AD) through a prism. Figure adapted from [6].

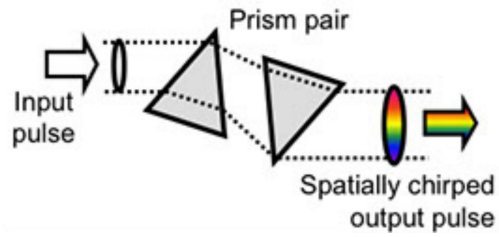

(b) Spatial Chirp (SC) induced by angular dispersion (AD) through a prism pair. Figure adapted from [7].

1.4.1 Theory of first-order STC

By *first-order* couplings, we refer to all terms that are proportional to the first power of the product xt. Another important point, as mentioned earlier, is that the spectral and temporal representations of the laser pulse are related by a Fourier transform. The same relationship holds between the spatial and spatial-frequency domains. Therefore, the electric field of a laser pulse can be equivalently expressed in four different spaces: (x,t), (x,ω), (k,t), and (k,ω). For a Gaussian laser pulse with STC, these representations take the following forms:

$$\begin{cases} E(x,t) \propto \exp\left[Q_{xx}x^2 + 2Q_{xt}xt - Q_{tt}t^2\right], \\ E(x,\omega) \propto \exp\left[R_{xx}x^2 + 2R_{x\omega}x\omega - R_{\omega\omega}\omega^2\right], \\ E(k,t) \propto \exp\left[P_{kk}k^2 + 2P_{kt}kt - P_{tt}t^2\right], \\ E(k,\omega) \propto \exp\left[S_{kk}k^2 + 2S_{k\omega}k\omega - S_{\omega\omega}\omega^2\right]. \end{cases} \tag{1.43}$$

Here, each coefficient is a complex number, with its real and imaginary parts representing different physical properties of the laser beam. For example:

$$Q_{xx} = -\frac{1}{w^2(z)} - \frac{i\pi}{\lambda R(z)}, \qquad Q_{tt} = \frac{1}{\tau^2} - i\beta.$$

Thus, $\mathrm{Re}\{Q_{xx}\}$ corresponds to the beam spot size (BSS), while $\mathrm{Im}\{Q_{xx}\}$ represents the wavefront curvature (WFC). Similarly, $\mathrm{Re}\{Q_{tt}\}$ describes the pulse duration (PD), whereas $\mathrm{Im}\{Q_{tt}\}$ corresponds to the temporal chirp (TCH). By extension, each of the twelve coefficients can be assigned a physical interpretation. Table 1 lists the cross terms responsible for STC.

Cross term	Real part (amplitude)	Imaginary part (phase)
Q_{xt}	Pulse-front tilt (PFT)	Wavefront rotation (WFR)
$R_{x\omega}$	Spatial chirp (SC)	Wavefront-tilt dispersion (WFD)
P_{kt}	Time vs angle (TVA)	Angular temporal chirp (ATC)
$S_{k\omega}$	Angular dispersion (AD)	Angular spectral chirp (ASC)

Table 1.1: Cross-term coefficients and their physical interpretations. Rows list the cross terms in each domain; columns indicate the meaning of their real and imaginary parts.

We already discussed earlier three of the well known amplitude couplings (AD,SC,PFT).The last one, time versus angle (TVA), causes the leading edge of the pulse to propagate at a different angle than the trailing edge. Wavefront rotation (WFR) produces a temporal rotation of the wavefront. Wavefront-tilt dispersion (WFD) introduces a frequency-dependent wavefront tilt. Angular temporal chirp (ATC) causes different propagation angles to accumulate distinct phase shifts over time, so that each angle experiences a varying instantaneous frequency. Angular spectral chirp (ASC) causes different propagation angles to accumulate distinct phase shifts across the spectrum, leading each angle to experience a varying group delay.

An important thing to keep in mind is that this 8 parameters (Real and imaginary) are not independent since again all of the four spaces are related via Fourier transformation. Which means that to fully characterise the coupled Gaussian beam (neglecting spatial aberrations) we only need to know six quantity. For example if we consider the spatio-temporal space (x,t) we need to know the BSS,WFC,PD,TCH,PFT and WFR. Another great

advantage of this formalism is that it allows to express explicit relation between the different couplings (cite).

$$\frac{\mathrm{Re}\{Q_{xt}\}}{\mathrm{Re}\{Q_{tt}\}} = -2\ \mathrm{Im}\{R_{x\omega}\} + 2\ \mathrm{Im}\{R_{\omega\omega}\}\,\frac{\mathrm{Re}\{R_{x\omega}\}}{\mathrm{Re}\{R_{\omega\omega}\}}. \tag{1.44}$$

Here the PFT term is normalised by the pulse duration and SC term is normalised by the pulse bandwidth. This relation captures the fact that pulse front tilt is the result of wavefront tilt dispersion and a combination of spatial chirp and frequency chirp. We can write similar relations for every coupled parameters:

$$\frac{\mathrm{Re}\{S_{k\omega}\}}{\mathrm{Re}\{S_{kk}\}} = -2\ \mathrm{Im}\{R_{x\omega}\} + 2\ \mathrm{Im}\{R_{xx}\}\,\frac{\mathrm{Re}\{R_{x\omega}\}}{\mathrm{Re}\{R_{xx}\}}.$$

Here the AD is normalised by angular divergence (AGD) and SC is normalised by the beam spot size. We can see that in the case of very large wavefront curvature,for example a collimated beam, the AD is proportional to WFD. In this approximation we can thus conclude that PFT is indeed caused by AD.

$$\frac{\mathrm{Re}\{R_{x\omega}\}}{\mathrm{Re}\{R_{xx}\}} = -2\ \mathrm{Im}\{S_{k\omega}\} + 2\ \mathrm{Im}\{S_{kk}\}\,\frac{\mathrm{Re}\{S_{k\omega}\}}{\mathrm{Re}\{S_{kk}\}}.$$

This relation shows that SC is caused by ASC and a combination of angular wavefront curvature and AGD.

$$\frac{\mathrm{Re}\{P_{kt}\}}{\mathrm{Re}\{P_{kk}\}} = -2\ \mathrm{Im}\{Q_{xt}\} + 2\ \mathrm{Im}\{Q_{xx}\}\,\frac{\mathrm{Re}\{Q_{xt}\}}{\mathrm{Re}\{Q_{xx}\}}.$$

The final relation shows that TVA is caused by WFR and a combination of wavefront curvature and PFT.

This theoretical framework allows to clarify the relationship between the different well known couplings. Furthermore it can be extended to more complex profiles using a numerical approach.

1.5 Plasma parameters

Plasma is a state of mater which can be described as an ionised medium where charged particles can move freely and where collective behaviours occur under the influence of external perturbations. These collective behaviours give rise to several characteristic quantities (such as characteristic length or response time). In this section we will describe some of these quantities that are important to understand the laser plasma interaction.

1.5.1 Electronic density

The electronic density describes the number of electrons per unit of plasma volume. In an unperturbed, fully ionised plasma, the ionic density n_i is constant, and the corresponding electronic density is $n_e = Zn_i$, where Z is the atomic number. At finite temperature, the electrons possess a velocity distribution due to their random thermal motion. In thermodynamic equilibrium, this velocity distribution follows the Maxwell–Boltzmann law,

$$f(v) \propto \exp\left(-\frac{mv^2}{2k_B T_e}\right), \tag{1.45}$$

where T_e is the electron temperature, k_B the Boltzmann constant, m the mass of the electron and v the velocity.

1.5.2 Plasma frequency

The plasma frequency characterise the natural frequency of the plasma response when submitted to an external perturbation. Let's consider a uniform unperturbed plasma of density n_e on which we displace a layer by a small distance δ_x. This displacement will result in two surfaces with surface charge $\sigma = \pm e n_e \delta_x$. Using Gauss law we find the resulting electric field between the two surfaces $E = \frac{\sigma}{\epsilon_0}$. The equation of motion is then:

$$m\frac{d^2\delta_x}{dt^2} = \frac{-e^2 n_e \delta_x}{\epsilon_0}$$

We recognise the equation of an harmonic oscillator with frequency:

$$\omega_p = \sqrt{\frac{e^2 n_e}{m\epsilon_0}} \tag{1.46}$$

The plasma frequency can also be viewed as the characteristic timescale of the plasma response with $\tau_p = \omega_p^{-1}$.

1.5.3 Critical density

The critical density $n_c = \frac{m\epsilon_0}{e^2}\omega_0^2$ is a threshold over which an electromagnetic wave of frequency ω_0 cannot propagate inside the plasma. We can illustrate this by considering the index of refraction of a cold unmagnetized plasma:

$$\eta = \sqrt{1 - \frac{\omega_p^2}{\omega_0^2}} = \sqrt{1 - \frac{n_e}{n_c}}$$

Where $n_e < n_c$ the refractive index is positive and the wave can propagate through the plasma.

When $n_e = n_c$ the refractive index becomes zero, which means the wave is evanescent and cannot propagate further.

For $n_e < n_c$, the refractive index becomes imaginary, causing the wave to decay exponentially. Effectively the plasma reflects the electromagnetic wave.

1.6 Ionisation processes

Ionisation process is necessary to produce a plasma from a gas target. Multiple processes exist to do so. In laser plasma interaction the main processes responsible for ionisation are barrier suppression ionisation (BSI), tunnel ionisation and multi photon ionisation. The rate at which each process is involved depends on the laser intensity. To illustrate BSI, let's consider an atomic potential in the presence of a static electric field E, the combined potential reads :

$$V(r) = -\frac{kZe^2}{r} - eEr \tag{1.47}$$

With $k = 4\pi/\epsilon_0$, Z the atomic number, e the electron charge and r the radial distance. To know the value of the electric field for an electron to escape the atomic potential we look for the position of the maximum of the total potential:

$$\frac{dV}{dr} = 0 \Rightarrow r_b = \sqrt{\frac{kZe}{E}} \Rightarrow V(r_b) = -2\sqrt{kZe^3E}$$

In order for the electron to become free of the nucleus we need $V(r_b) \leq -W_i$ where W_i is the ionisation potential of the atom. This leads to the electric field for BSI:

$$E_{BSI} = \frac{W_i^2}{2kZe^3} \Rightarrow I_{BSI} = \frac{1}{2}\epsilon_0 c E_{BSI}^2 [W/cm^2] \tag{1.48}$$

To give an example, $I_{BSI} = 1.37 \times 10^{14} W/cm^2$ for H^+. It is 8.77×10^{15} for He^{2+}, 1.47×10^{16} for N^{5+} and 1.61×10^{19} for N^{7+}.

Tunnel ionisation is a quantum process that occurs at lower intensities when the atomic potential barrier is distorted enough by the laser field such that an electron can tunnel through it and be freed even if $V(r) \geq -W_i$. On the other hand, multi photon ionisation is a process by which multiple photons with energy lower than the ionisation potential combine in order to ionise an atom. The Keldysh parameter characterizes which of these processes dominates

$$\gamma = \frac{\omega_0 \sqrt{W_i}}{E_0}$$

With ω_0 the laser frequency and E_0 the electric field amplitude. When $\gamma \gg 1$ multiphon ionisation dominates whereas when $\gamma \lesssim 1$ tunnel ionisation dominates. In the

regimes that interest us, $I \sim 10^{18}[W/cm^2]$ and $\omega_0 = 0.8\mu m$ tunnel ionisation or BSI dominate depending on the gas. In FIG. 1.9 we present an illustration of the BSI and tunneling ionisation process.

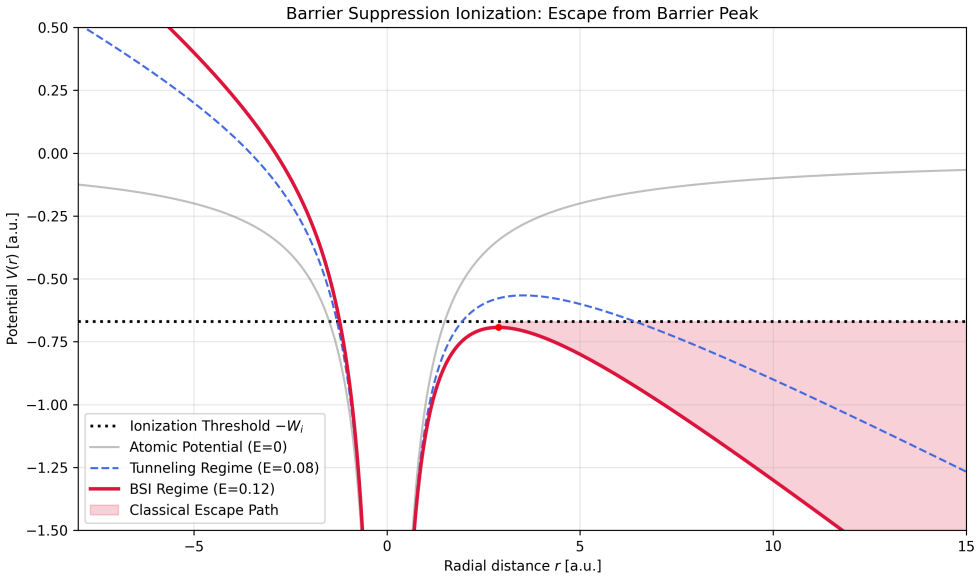

Figure 1.9: Comparison of field-induced ionisation regimes. The plot illustrates the deformation of the atomic potential $V(r)$ under different electric field strengths. The solid gray line represents the undisturbed symmetric Coulomb potential ($E = 0$). The dashed blue line represents the tunneling regime ($E = 0.08$ a.u.), where the potential barrier is distorted but its maximum remains above the ionisation threshold $-W_i$ (indicated by the horizontal black dotted line). The solid red line illustrates the Barrier Suppression Ionisation (BSI) regime ($E = 0.12$ a.u.), where the barrier is suppressed below the ionisation level. The shaded red area highlights the classical escape path for the electron. On the left side ($r < 0$), the potential is raised due to the $+eEr$ contribution, while on the right side ($r > 0$), the $-eEr$ term lowers the barrier, enabling ionisation.

1.7 Electrons in an electromagnetic wave

We will start by considering the Lagrangian of a charged particle in an Electromagnetic field.

$$\mathcal{L} = -mc^2\sqrt{1 - \frac{v^2}{c^2}} - q\phi + q\mathbf{v}.\mathbf{A} \tag{1.49}$$

With c the speed of light, v the velocity of the charged particle, q the charge ($-e$ for an electron), ϕ the scalar potential (null here since we consider a pure EM wave) and $\mathbf{A}$

the vector potential. From this Lagrangian the canonical momentum and Hamiltonian yield:

$$\begin{cases} \mathbf{P} = \mathbf{p} - e\mathbf{A} \\ \mathcal{H} = \gamma mc^2 \end{cases}$$

With $\mathbf{p} = \gamma m\mathbf{v}$ the momentum and $\gamma = \sqrt{1 + \frac{\mathbf{p}^2}{m^2c^2}}$ the Lorentz factor. Now if we consider a plane wave propagating along the z the direction and linearly polarised in the (x,y) plane such that $\mathbf{A}(r,t) \equiv \mathbf{A}(z,t)$ then $\frac{\partial \mathcal{L}}{\partial r_\perp} = 0$, i.e $\mathbf{P}_\perp$ is a constant of motion. With this we can derive the equations of motion:

$$\frac{d}{dt}(\mathbf{p}_\perp - e\mathbf{A}) = 0$$

$$\frac{d}{dt}p_z = -e\mathbf{v}_\perp . \frac{\partial \mathbf{A}}{\partial z}$$

We recall the work W done by the particle:

$$W = \int -e\mathbf{v}.\mathbf{E}dt \tag{1.50}$$

With $\mathbf{E} = -\frac{\partial \mathbf{A}}{\partial t}$ the electric field. For an electron initially at rest $\mathbf{v}_\perp = \frac{e\mathbf{A}}{m\gamma}$, the equation of energy conservation is:

$$\frac{d}{dt}\mathcal{H} = \frac{e^2}{2m\gamma}\frac{\partial A^2}{\partial t} \tag{1.51}$$

using the conservation of the canonical momentum, the z component of the equations of motion becomes:

$$\frac{d}{dt}p_z = \frac{-e^2}{2m\gamma}\frac{\partial A^2}{\partial z} \tag{1.52}$$

Therefore:

$$\frac{d}{dt}(\gamma mc^2 - p_z c) = \frac{e^2}{2m\gamma}\left(\frac{\partial}{\partial t} + c\frac{\partial}{\partial z}\right)A^2 \tag{1.53}$$

Now we introduce the phase variable $\xi = t - \frac{z}{c}$ that follows the wavefront of the EM wave , $\frac{\partial}{\partial t} = \frac{\partial}{\partial \xi}$, $\frac{\partial}{\partial z} = \frac{-1}{c}\frac{\partial}{\partial \xi}$. In this coordinate system the later equation reduces to:

$$\frac{d}{dt}(\gamma mc^2 - p_z c) = 0 \tag{1.54}$$

So for a particle initially at rest we have:

$$\gamma = 1 + \frac{p_z}{mc} = \sqrt{1 + \frac{\mathbf{p}^2}{m^2c^2}}$$

Finally we get:

$$\begin{cases} \frac{p_z}{mc} = \frac{1}{2}\left(\frac{e\mathbf{A}}{mc}\right)^2 \\ \gamma = 1 + \frac{1}{2}\left(\frac{e\mathbf{A}}{mc}\right)^2 \end{cases}$$

We define the normalized vector potential as

$$a_0 = \frac{e\,|\mathbf{A}|}{mc}. \tag{1.55}$$

This parameter is a key characteristic in laser–plasma interaction, as it determines the interaction regime. For $a_0 \gg 1$, the motion becomes ultrarelativistic and predominantly longitudinal. In contrast, the classical regime is identified by $a_0 \ll 1$, where the electron motion is mainly transverse because the magnetic component of the Lorentz force can be neglected.

In fact, one may characterize the electron motion by the angle its momentum makes with the propagation axis,

$$\frac{p_\perp}{p_z} = \sqrt{\frac{2}{\gamma - 1}} = \tan\theta, \tag{1.56}$$

which shows that the angle decreases as γ increases. In laser–plasma acceleration, we are interested in the regime where $a_0 \geq 1$.

In practical units, a_0 is often expressed as:

$$a_0 = 0.85\sqrt{I\,[10^{18}\,\mathrm{W/cm^2}]\,\lambda^2\,[\mu\mathrm{m}^2]}. \tag{1.57}$$

This expression enables direct estimation of the normalized vector potential from the laser intensity I and wavelength λ.

1.8 Fluid description of laser plasma interaction

In this section we are going to establish the fluid equations used to describe the laser plasma interaction. We are first going to make a couple of hypothesis. The plasma is considered cold, i.e we neglect the thermal velocity of the electrons, this is justified by the fact that the thermal velocity is much smaller than the one induced by the laser. We will consider only an electron fluid and assume the ions are immobile. This assumption is made since we have $\omega_{pe} \gg \omega_{pi}$: the response time of the ions is very long in comparison with the one of the electrons. We will use the following normalisation in what follows: $t \equiv \omega_0 t$, $\mathbf{r} \equiv \frac{\omega_0}{c}\mathbf{r}$, $\mathbf{v} \equiv \frac{\mathbf{v}}{c}$, $\mathbf{p} \equiv \frac{\mathbf{p}}{mc}$, $\mathbf{A} \equiv \frac{e\mathbf{A}}{mc}$, $\phi \equiv \frac{e\phi}{mc^2}$, $n \equiv \frac{n}{n_c}$ the electron density, where $n_c = \frac{m\epsilon_0\omega_0^2}{e^2}$ is the critical density.

Using these definitions the equation of motion is:

$$\left(\frac{\partial}{\partial t} + \mathbf{v}.\nabla\right)\mathbf{p} = \frac{\partial \mathbf{A}}{\partial t} + \nabla\phi - \mathbf{v} \times \nabla \times \mathbf{A} \tag{1.58}$$

using the following vectorial identity:

$$\nabla(\mathbf{A}\cdot\mathbf{B}) = (\mathbf{A}\cdot\nabla)\mathbf{B} + (\mathbf{B}\cdot\nabla)\mathbf{A} + \mathbf{A}\times(\nabla\times\mathbf{B}) + \mathbf{B}\times(\nabla\times\mathbf{A})$$

we find :

$$(\mathbf{v}.\nabla)\mathbf{p} = \nabla\gamma - \mathbf{v} \times \nabla \times \mathbf{p}$$

Therefore the equation of motion becomes:

$$\frac{\partial}{\partial t}(\mathbf{p} - \mathbf{A}) = \nabla(\phi - \gamma) + \mathbf{v} \times \nabla \times (\mathbf{p} - \mathbf{A}) \tag{1.59}$$

Taking the curl of the last equation we get:

$$\frac{\partial}{\partial t}\nabla \times (\mathbf{p} - \mathbf{A}) = \nabla \times \Big[\mathbf{v} \times \nabla \times (\mathbf{p} - \mathbf{A})\Big]. \tag{1.60}$$

At $t = 0$ we have $\mathbf{p} = 0$, $\mathbf{A} = 0$, $\nabla \times (\mathbf{p} - \mathbf{A}) = 0$. Since the evolution equation is homogeneous for $\nabla \times (\mathbf{p} - \mathbf{A})$, the uniqueness theorem implies that $\nabla \times (\mathbf{p} - \mathbf{A})$ is identically zero.

Finally the equation of motion:

$$\frac{\partial}{\partial t}(\mathbf{p} - \mathbf{A}) = \nabla(\phi - \gamma) \tag{1.61}$$

In this equation, two source terms can be identified for the evolution of the canonical momentum. The term $\nabla\phi$ represents the Coulomb (electrostatic) force, which tends to push electrons toward regions of lower density. The term $-\nabla\gamma$ corresponds to the ponderomotive force, which drives electrons toward regions of lower field intensity.

Since there is four terms describing the fluid we need four equations for the system to be self consistent. The next equation is the continuity equations which expresses the conservation of the charge by equating the particle flux with the variation of the density per unit of time:

$$\frac{\partial n}{\partial t} + \nabla.\left(\frac{n\mathbf{p}}{\gamma}\right) = 0 \tag{1.62}$$

Next we will use the Maxwell Ampere equation with the coulomb Gauge ($\nabla.\mathbf{A} = 0$) to express the evolutions of the potentials under the influence of plasma current.

$$\left(\frac{\partial^2}{\partial t^2} - \Delta\right)\mathbf{A} + \frac{\partial\nabla\phi}{\partial t} = -\frac{n\mathbf{p}}{\gamma} \tag{1.63}$$

Finally we will use the Poisson equation to express the evolution of the scalar potential under the influence of density variations. For this we will express the charge density as follows:

$$\rho = -e(n - n_0)$$

whith $n_0 = Zn_i$ the unperturbed electron density, n_i the ion density. In what follows $n_0 \equiv \frac{n_0}{n_c}$. Using Maxwell Gauss equation we can then express the Poisson equation:

$$\Delta\phi = \frac{\omega_p^2}{\omega_0^2}\left(\frac{n}{n_0} - 1\right) \tag{1.64}$$

We arrived at a self-consistent system of coupled non linear partial differential equations, that captures the dynamic of the laser plasma interaction under the hypothesis made at the begin of this section.

$$\begin{cases} \frac{\partial}{\partial t}(\mathbf{p} - \mathbf{A}) = \nabla(\phi - \gamma) \\ \frac{\partial n}{\partial t} + \nabla \cdot \left(\frac{n\mathbf{p}}{\gamma}\right) = 0 \\ \left(\frac{\partial^2}{\partial t^2} - \Delta\right)\mathbf{A} + \frac{\partial \nabla\phi}{\partial t} = -\frac{n\mathbf{p}}{\gamma} \\ \Delta\phi = \frac{\omega_p^2}{\omega_0^2}\left(\frac{n}{n_0} - 1\right) \end{cases}$$

1.8.1 Linear solution

To gain an initial understanding of the system dynamics, we consider the linear regime ($a_0 \ll 1$). In this case, the small value of a_0 allows us to linearize the equations using a perturbative approach. Specifically, we assume that the deviations from equilibrium of the four quantities describing the system are small, such that $|\mathbf{p}| \ll 1$, $n = n_0 + \delta n$ with $\delta n \ll n_0$, and $\gamma \simeq 1 + \frac{\mathbf{p}^2}{2}$.

Using this ansatz, the continuity and Poisson equations become:

$$\frac{\partial}{\partial t}\frac{\delta n}{n_0} + \nabla \cdot \mathbf{p} = 0 \tag{1.65}$$

$$\Delta\phi = \frac{\omega_p^2}{\omega_0^2}\frac{\delta n}{n_0} \tag{1.66}$$

To solve this system in the linear regime, we make a key assumption. We consider an underdense plasma, $\omega_p \ll \omega_0$. This allows us to decompose the plasma response into fast and slow components, i.e.

$$\mathbf{p} = \mathbf{p}_{\text{fast}}(\omega_0 t) + \mathbf{p}_{\text{slow}}(\omega_p t).$$

Substituting this decomposition into the transverse component of the equation of motion gives:

$$\frac{\partial\left(\mathbf{p}_{\perp,\text{slow}} + \mathbf{p}_{\perp,\text{fast}} - \mathbf{A}_\perp\right)}{\partial t} = \nabla_\perp(\phi - \gamma) \tag{1.67}$$

The right-hand side of this equation drives the slow plasma motion. Since the transverse laser beam size is on the order of the plasma wavelength or larger, we can approximate $\nabla_\perp \sim 1/L_\perp$, with $L_\perp$ the normalised transverse beam size. The equation then becomes:

$$p_{\perp,\text{fast}} - A_\perp \sim \frac{1}{L_\perp}\left((\phi-\gamma) - p_{\perp,\text{slow}}\right) \lesssim \frac{\omega_p}{\omega_0}\left((\phi-\gamma) - p_{\perp,\text{slow}}\right). \tag{1.68}$$

From this ordering we may neglect the right-hand side. Additionally using the gauge condition and the same assumption leads to $A_\perp \gg A_z$ and

$$\mathbf{p}_{\perp,\text{fast}} \simeq \mathbf{A}_\perp \simeq \mathbf{A}.$$

Next, consider the longitudinal component of the equation of motion, assuming that the laser pulse length is comparable to the plasma wavelength:

$$p_{z,\text{fast}} \sim \frac{\omega_p}{\omega_0}\left((\phi-\gamma) - p_{z,\text{slow}}\right). \tag{1.69}$$

This shows that:

$$\mathbf{p}_{z,\text{fast}} \ll \mathbf{p}_{\perp,\text{fast}} \simeq \mathbf{p}_{\text{fast}}.$$

Returning to the full equation of motion, we obtain:

$$\frac{\partial \mathbf{p}_{\text{slow}}}{\partial t} = \nabla\left(\phi - \frac{\mathbf{p}_{\text{slow}}^2}{2} - \frac{\mathbf{p}_{\text{fast}}^2}{2} - \mathbf{p}_{\text{slow}}\cdot\mathbf{p}_{\text{fast}}\right). \tag{1.70}$$

This expression simplifies upon time averaging over one optical cycle. In this case, all first-order terms in the fast motion vanish: $\langle\mathbf{A}\rangle = 0$, $\langle\mathbf{p}\rangle = \mathbf{p}_{\text{slow}}$. Thus,

$$\frac{\partial \mathbf{p}_{\text{slow}}}{\partial t} = \nabla\left(\phi - \frac{\mathbf{p}_{\text{slow}}^2}{2} - \frac{\langle\mathbf{A}^2\rangle}{2}\right). \tag{1.71}$$

As discussed previously, for $a_0 \ll 1$ the electron motion is predominantly transverse inside the laser pulse (fast quiver motion). Consequently, $\mathbf{p}_{\text{slow}}^2$ does not have time to build up significantly over one optical period and is therefore negligible compared to $\langle\mathbf{A}^2\rangle$. This yields:

$$\frac{\partial \mathbf{p}_{\text{slow}}}{\partial t} = \nabla\left(\phi - \frac{\langle\mathbf{A}^2\rangle}{2}\right). \tag{1.72}$$

Combining the last equations we arrive at:

$$\Delta\left(\frac{\partial^2}{\partial t^2} + \frac{\omega_p^2}{\omega_0^2}\right)\phi = \frac{\omega_p^2}{\omega_0^2}\Delta\frac{\langle\mathbf{A}^2\rangle}{2} \tag{1.73}$$

By integration this yields:

$$\left(\frac{\partial^2}{\partial t^2}+\frac{\omega_p^2}{\omega_0^2}\right)\phi=\frac{\omega_p^2}{\omega_0^2}\frac{\langle\mathbf{A}^2\rangle}{2} \tag{1.74}$$

For convenience we will use the new coordinate system (τ,ξ) with $\tau=t$ and $\xi=z-t$ (in normalized unit). Using this coordinate system that follows the laser and the quasistatic approximation which neglects the temporal derivative, the laser appears frozen on the timescale of the interaction ($\omega_p \gg c/z_r$) and the last equation becomes:

$$\left(\frac{\partial^2}{\partial \xi^2}+\frac{\omega_p^2}{\omega_0^2}\right)\phi=\frac{\omega_p^2}{\omega_0^2}\frac{\langle\mathbf{A}^2\rangle}{2} \tag{1.75}$$

The laser is considered Gaussian with its envelope:

$$a^2(r,\xi)=a_0^2\exp\left(\frac{-\xi^2}{L_0^2}\right)\exp\left(\frac{-r^2}{\sigma^2}\right)$$

With L_0 the normalised pulse length and σ, the normalised transverse beam size.
We thus have:

$$\langle\mathbf{A}^2\rangle=\frac{1}{T_0}\int_{-T_0/2}^{T_0/2}a^2\cos^2(\xi)\,dt\;\simeq\frac{a^2}{2} \tag{1.76}$$

To solve this integral we use again the quasistatic approximation assuming that a^2 does not depend on time. This is justified by the fact that a^2 is almost constant on one optical cycle. Also in this integral z is considered constant therefore $dt=d\xi$. The final equation on the scalar potential becomes:

$$\left(\frac{\partial^2}{\partial \xi^2}+\frac{\omega_p^2}{\omega_0^2}\right)\phi=\frac{\omega_p^2}{\omega_0^2}\frac{a^2}{4} \tag{1.77}$$

This equation is a standard second order ordinary differential equation with a source term. It can be solved by finding the appropriate Green function. For that we know that $\phi=0$ in front of the laser therefore we have $\phi(\xi\to+\infty)=0$. This leads to the following Green function:

$$G(\xi,\xi')=\begin{cases}0 & \text{if } \xi>\xi'\\ -\frac{\omega_0}{\omega_p}\sin\left(\frac{\omega_p}{\omega_0}(\xi-\xi')\right) & \text{if } \xi<\xi'\end{cases} \tag{1.78}$$

Therefore the potential reads:

$$\phi(r,\xi)=-\frac{\omega_p}{4\omega_0}\int_{\xi}^{+\infty}a^2\sin\left(\frac{\omega_p}{\omega_0}(\xi-\xi')\right)\,d\xi' \tag{1.79}$$

Since we are interested in the perturbation behind the laser pulse, i.e $\xi < 0$ and a^2 is centered around $\xi = 0$, rapidly vanishes for $\xi < 0$, the integral can be written from $-\infty$ to $+\infty$

$$\phi(r,\xi) = -\frac{\omega_p}{4\omega_0}\left[\sin\left(\frac{\omega_p}{\omega_0}(\xi)\right)\int_{-\infty}^{+\infty} a^2\cos\left(\frac{\omega_p}{\omega_0}(\xi')\right)\,d\xi' \;-\cos\left(\frac{\omega_p}{\omega_0}(\xi)\right)\int_{-\infty}^{+\infty} a^2\sin\left(\frac{\omega_p}{\omega_0}(\xi')\right)\,d\xi'\right] \tag{1.80}$$

The second integral in the RHS is zero since a^2 is an even function. The first integral is thus equivalent to the Fourier transform of a^2 with conjugate variable $\frac{\omega_p}{\omega_0}$

Finally we find:

$$\phi(r,\xi) = -\frac{\sqrt{\pi}\omega_p L_0 a_0^2}{4\omega_0}\exp\left(\frac{-\omega_p^2 L_0^2}{4\omega_0^2}\right)\exp\left(\frac{-r^2}{\sigma^2}\right)\sin\left(\frac{\omega_p}{\omega_0}\xi\right) \tag{1.81}$$

We can therefore find the electric field using $\mathbf{E} = -\nabla\phi$

$$E_r = -\frac{\sqrt{\pi}\omega_p L_0 a_0^2 r}{2\sigma^2\omega_0}\exp\left(\frac{-\omega_p^2 L_0^2}{4\omega_0^2}\right)\exp\left(\frac{-r^2}{\sigma^2}\right)\sin\left(\frac{\omega_p}{\omega_0}\xi\right) \tag{1.82}$$

$$E_z = \frac{\sqrt{\pi}\omega_p^2 L_0 a_0^2}{4\omega_0^2}\exp\left(\frac{-\omega_p^2 L_0^2}{4\omega_0^2}\right)\exp\left(\frac{-r^2}{\sigma^2}\right)\cos\left(\frac{\omega_p}{\omega_0}\xi\right) \tag{1.83}$$

As for the density perturbation we can find it using the Poisson equation:

$$\frac{\delta n}{n_0} = \frac{\sqrt{\pi}\omega_p L_0 a_0^2}{4\omega_0}\left[1 + \frac{2\omega_0^2}{\omega_p^2\sigma^2}\left(1 - \frac{2r^2}{\sigma^2}\right)\right]\exp\left(\frac{-\omega_p^2 L_0^2}{4\omega_0^2}\right)\exp\left(\frac{-r^2}{\sigma^2}\right)\sin\left(\frac{\omega_p}{\omega_0}\xi\right) \tag{1.84}$$

The linear solution (see FIG. 1.10) shows that the plasma wave and the associated longitudinal and transverse electric fields resulting from the interaction with a laser pulse with $a_0 \ll 1$ is sinusoidal. Its amplitude decays with the radial coordinate. The wavelength of this plasma wave is $\lambda_p = 2\pi/k_p$. The maximum value of the longitudinal electric field is found on axis and is :

$$E_{zmax} = \frac{E_0\sqrt{\pi}\omega_p L a_0^2}{4c}$$

with $E_0 = mc\omega_p/e$ the cold non relativistic wavebreaking field which represents the maximum value of the longitudinal electric field the plasma wave can sustain before the model breaksdowns due to a singularity of the density. L is the pulse length. An important fact to note is that the energy transfer between the laser pulse and the plasma wave is maximised for a pulse length on the order of the plasma wavelength. Indeed we have for $\frac{dE_z}{dL} = 0$, $L = \frac{\sqrt{2}}{k_p}$. In addition to the longitudinal field there is a radial electric field which is zero on axis but acts either as a focusing or defocusing element depending

on the position relative to the laser pulse. The maximum of this radial field is found for $r = \frac{\sigma}{\sqrt{2}}$.

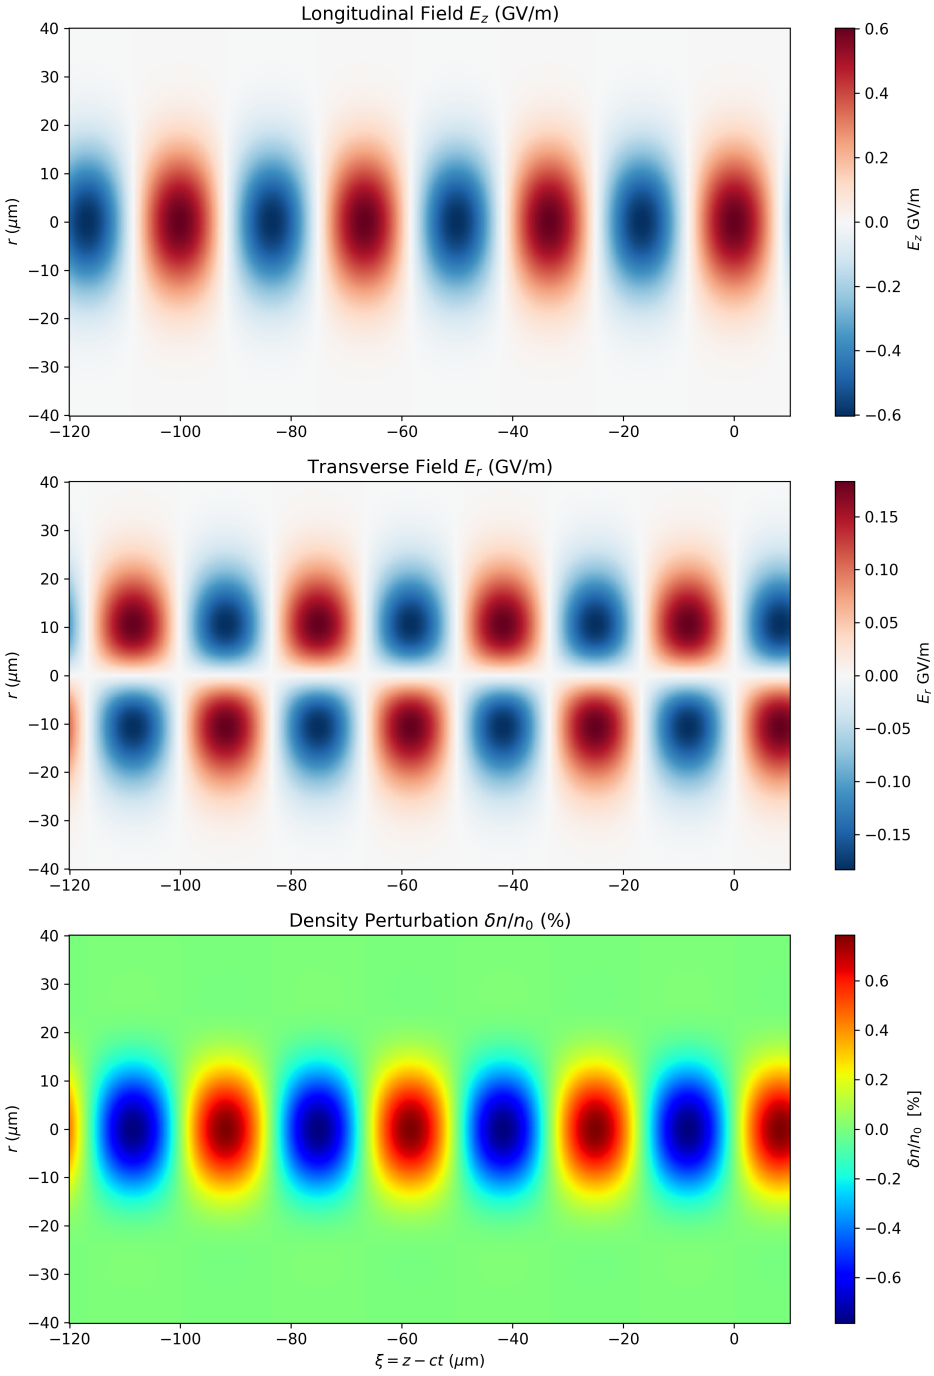

Figure 1.10: Linear solution for $a_0 = 0.1$, $\lambda = 800\,\mathrm{nm}$, $w_0 = 15\,\mu\mathrm{m}$,$n_0 = 10^{18} cm^{-3}$ and $L = \sqrt{2}/k_p$. The top panel represents the longitudinal electric field, the middle panel the transverse electric field, and the bottom panel the plasma density perturbation.

1.8.2 1D Nonlinear plasma waves

To study the nonlinear dynamic we use the same assumption as the one of the previous sections but with this time $a_0 > 1$. Furthermore we will consider a 1D model where all the quantities only depend on z. This assumption is valid for transversely broad pulse i.e, $k_p r \gg 1$. Using the quasistatic approximation and the phase variable $\xi = z - t$ the equation of motion along the z direction:

$$\frac{\partial}{\partial \xi}(\phi - \gamma(1 - \beta_z)) = 0 \tag{1.85}$$

To establish this equation we used $p_z = \frac{\gamma v_z}{c} = \gamma \beta_z$. The continuity equation:

$$\frac{\partial}{\partial \xi}(n(1 - \beta_z)) = 0 \tag{1.86}$$

Using for initial condition $v_z = 0$,$\phi = 0$,$n = n_0$ we find:

$$\gamma(1 - \beta_z) = 1 + \phi \tag{1.87}$$

$$n(1 - \beta_z) = n_0 \tag{1.88}$$

Now using $\gamma^2 = 1 + \mathbf{A}^2 + p_z^2$ we can write for the motion equation:

$$\frac{\sqrt{1 + \mathbf{A}^2}}{\sqrt{1 - \beta_z^2}}(1 - \beta_z) = 1 + \phi \tag{1.89}$$

which leads after some algebraic manipulation:

$$\beta_z = \frac{(1 + \mathbf{A}^2) - (1 + \phi^2)}{(1 + \mathbf{A}^2) + (1 + \phi^2)} \tag{1.90}$$

Using the continuity equation we have:

$$\frac{n}{n_0} - 1 = \frac{\beta_z}{(1 - \beta_z)} \tag{1.91}$$

Now by injecting the last expression into (1.64) we find the wake equation:

$$\frac{\partial^2}{\partial \xi^2}\phi = \frac{\omega_p^2}{2\omega_0^2}\left[\frac{1 + \mathbf{A}^2}{(1 + \phi)^2} - 1\right] \tag{1.92}$$

For the quasi-static approximation to hold in equation (1.92), the group velocity of the laser must be close to the speed of light, $v_g \simeq c$. However, if v_g significantly differs from c, we can still use the quasi-static approximation by shifting to a moving frame that follows the plasma wavefront. In typical laser-plasma conditions, the group velocity of

the laser is approximately equal to the phase velocity of the plasma wake, $v_g \simeq v_\phi$, therefore we introduce a new phase variable:

$$\xi = z - v_\phi t,$$

where the phase velocity is given by

$$v_\phi = \frac{\omega_0}{k_0 c} = \beta_p$$

Following the same procedure as before, but now in this moving frame, we obtain the modified wakefield equation:

$$\frac{\partial^2 \phi}{\partial \xi^2} = \frac{\omega_p^2}{\omega_0^2}\gamma_p^2 \left(\beta_p \left[1 - \frac{1 + \mathbf{A}^2}{\gamma_p^2 (1+\phi)^2} \right]^{-1/2} - 1 \right), \tag{1.93}$$

where $\gamma_p = (1 - \beta_p^2)^{-1/2}$ is the Lorentz factor associated with the phase velocity.

Analytical solutions of equations (1.93) exist for square laser pulse profile[8, 9]. For large a_0 these solutions describe a plasma wave that deviates from the simple sinusoidal solution. Indeed the periodic pattern steepens and its period gets longer. The potential oscillates between this minimum and maximum value for a linearly polarized laser with optimal length:

$$\phi_m = \hat{E}_{max}^2/2 \pm \beta_p \sqrt{(1 + \hat{E}_{max}^2/2)^2 - 1}$$

$$\hat{E}_{max} = \frac{(a_0^2/2)}{\sqrt{1 + (a_0^2/2)}}$$

Where $\hat{E}_{max} = E_{max}/E_0$. We can immediately see that for large value of a_0 the amplitude of the plasma wave is considerably stronger than in the linear case. As for the lengthening of the period we have for $a_0 \gg 1$:

$$\lambda_{Np} \simeq \frac{2\lambda_p a_0}{\pi}$$

In the ultra relativistic case this lengthening can be important and therefore considerably increase the accelerating length and potentially the energy gain.

1.9 Particle trapping

In order to accelerate electrons we need to be able to trap them inside the plasma wave generated by the laser. Several methods exist and they are all based on the same principle. In this section we will explain how this phenomenon occurs using the Hamiltonian

formalism. Let's consider the Lagrangian (1.49). Using this Lagrangian the Hamiltonian of an electron in 1D nonlinear plasma wave yields:

$$\mathcal{H}(z, p_z) = \gamma(p_z) - \phi(z) \tag{1.94}$$

We express this Hamiltonian in the frame of the plasma wavefront. Therefore we will make a canonical change of variable from $(z, p_z) \rightarrow (\xi, p_\xi)$. This leads to the introduction of the following generating function:

$$F(z, p_\xi, t) = (z - v_\phi t)p_\xi$$

which leads to:

$$\xi = \frac{\partial F}{\partial p_\xi} = z - v_\phi t \quad p_z = \frac{\partial F}{\partial z} = p_\xi$$

Finally we have:

$$\mathcal{H}(\xi, p_\xi) = \mathcal{H}(z, p_z) + \frac{\partial F}{\partial t} = \gamma - \beta_p p_z - \phi \tag{1.95}$$

An important property of this Hamiltonian is that it has no explicit time dependence and thus it is conserved along a particle trajectory in phase space. Using this fact and by considering an electron with initial energy H_0 we can solve for its longitudinal momentum:

$$p_z = \beta_p \gamma_p^2 (H_0 + \phi) \pm \gamma_p \sqrt{\gamma_p^2 (H_0 + \phi)^2 - (1 + \mathbf{A}^2)} \tag{1.96}$$

This equation allows us to draw the particle trajectories for a given initial energy, with ϕ the solution to equation (1.93). There is two special trajectories, the first one corresponds to electrons born initially at rest in front of the laser pulse. For these electrons, ie $H_0 = 1$, $p_z(\xi_i) = \mathbf{A}(\xi_i) \simeq 0$, where ξ_i indicates the phase at which the electron is ionised. These electron trajectories are referred to as the cold fluid orbit, they are not trapped in the plasma wave but they contribute to the formation of the plasma wakefield. The other special trajectory is called the separatrix which separates trapped from untrapped orbits. An example of the resulting phase space and wakefield structure is shown in Fig. 1.11.To find the separatrix we need to study the critical points of the Hamiltonian. These critical points are characterized by:

$$\frac{\partial \mathcal{H}}{\partial p_z} = 0, \quad \frac{\partial \mathcal{H}}{\partial \xi} = 0$$

which leads to (behind the laser pulse):

$$p_z = \gamma \beta_p, \quad \frac{\partial \phi}{\partial \xi} = 0$$

We find that stable critical points are located at $v_z = v_\phi$ and $\phi = \phi_{max}$, while the unstable critical points are located at $v_z = v_\phi$ and $\phi = \phi_{min}$. Therefore the separatrix is defined by the contour:

$$H_{sep} = H(p_z = \gamma\beta_p, \xi = \xi_{min}) = \frac{\sqrt{1 + \mathbf{A}^2(\xi_{min})}}{\gamma_p} - \phi_{min}$$

where ξ_{min} corresponds to $\phi_{min} = \phi(\xi_{min})$. To be trapped an electron orbit must satisfy the condition $H \leq H_{sep}$. On the plot of the phase space we can see that the width of the separatrix corresponds to the nonlinear plasma wavelength. This Hamiltonian formalism allows us to determine the maximum energy gain a trapped electron can experience. Since the Hamiltonian is conserved, the maximum energy gain is obtained by equating the Hamiltonian that corresponds to the maximum energy with the Hamiltonian of a particle just inside the separatrix,i.e $H(p_{zmax}, \xi_{max}) = H_{sep}$. Using equation (1.96) this yields:

$$p_{zm} = \beta_p\gamma_p^2(H_{sep} + \phi) \pm \gamma_p\sqrt{\gamma_p^2(H_{sep} + \phi)^2 - (1 + \mathbf{A}^2(\xi_{max}))} \tag{1.97}$$

which leads in the limit $\gamma_p^2 \gg 1$ and $\gamma_p\delta\phi \gg 1$, with $\delta\phi = \phi_{max} - \phi_{min}$:

$$\begin{cases} p_{z\text{max}} \simeq 2\gamma_p^2\delta\phi \\ p_{z\text{min}} \simeq \frac{1}{2\delta\phi} - \frac{\delta_\phi}{2} \\ \gamma_{\text{max}} \simeq 2\gamma_p^2\hat{E}_{max}^2 \quad if \quad \hat{E}_{max}^2 \gg 2 \\ \gamma_{\text{max}} \simeq 4\gamma_p^2\hat{E}_{max} \quad if \quad \hat{E}_{max}^2 \ll 2 \end{cases}$$

Once again we see that in the relativistic case the Energy gain is much stronger than in the linear case. Moreover with this formalism we can estimate the optimal length for maximum energy gain. This length is called the dephasing length and corresponds to the maximum distance a trapped electron can travel in the plasma wave before it encounters a decelerating longitudinal electric field. To evaluate this length we can just equate the maximum energy gain with the work done by the same electron, this yields in the case $a_0 \gg 1$:

$$mc^2\gamma_{max} = eE_{max}L_d$$

$$L_d = \frac{\gamma_p^2\lambda_{NP}}{2}$$

The expression for the dephasing length shows that it is proportional to the nonlinear plasma wavelength; thus, higher energy gain can potentially be achieved by increasing the laser amplitude a_0.

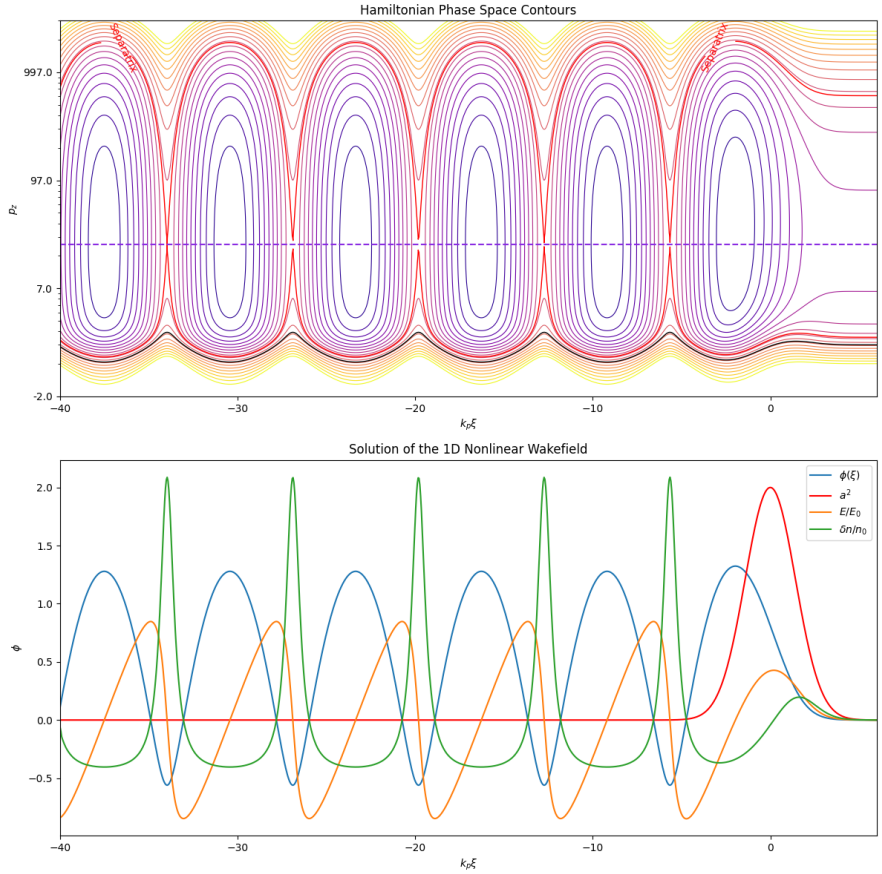

Figure 1.11: 1D nonlinear wakefield and phase-space solution of Eq. (1.93) for a Gaussian laser pulse with $a_0 = 2$, $\lambda = 800\,\mathrm{nm}$, $\gamma_p = 22.37$, $\omega_p/\omega_0 = 0.01$, and pulse width $L = 25\,\mu\mathrm{m}$. The top panel shows the phase space, with the Hamiltonian represented as contours. The red contour line represents the separatrix, while the black contour line corresponds to the cold-fluid orbit. The horizontal dashed purple line represents the wake phase velocity $v = v_\phi$. The bottom panel shows the associated wakefields, namely the squared vector potential, the scalar potential, the normalized longitudinal electric field, and the density perturbation.

1.10 Blowout regime

In the previous sections, I presented the linear regime, which is three-dimensional and assumes $a_0 \ll 1$. I also discussed the one-dimensional nonlinear regime, which assumes a transversely broad laser pulse ($k_p r \gg 1$). However, in realistic scenarii both assumptions break down. In typical LWFA experiments we usually have $a_0 \gtrsim 1$ and

$k_p r \lesssim 1$. In this regime, if the laser intensity is sufficiently high, all electrons near the optical axis are expelled. Moreover, the nonlinear plasma wavelength can exhibit a radial dependence, leading to a curved plasma wavefront. This regime is commonly referred to as the *blowout* (or *bubble*) regime, where the accelerating structure is an ion cavity almost completely devoid of electrons. This regime has been thoroughly studied using numerical simulations, and it presents many advantages for electron acceleration.

Firstly, as in the 1D nonlinear regime, the longitudinal electric field can reach very high values, on the order of 100 GV/m. For example, it has been shown [10] that for laser amplitudes $a_0 \gtrsim 2$, and for a matched laser spot size $r_0 \simeq r_b$, with r_0 and r_b denoting the laser spot size and the plasma bubble radius respectively, the peak accelerating field scales as $\hat{E}_{\text{max}} \simeq \sqrt{a_0}$.

In addition, the spherical shape of the ion cavity provides linear transverse focusing forces, i.e., $F_r \propto -r$. This is a crucial property, as it preserves the transverse emittance of the accelerated electron bunch.

Another key advantage is that the ion cavity can act as a *self-guiding* structure: it focuses the laser pulse and counteracts its natural diffraction. As a result, the acceleration length can extend beyond several Rayleigh lengths (see later section on self-focusing).

Finally, in this regime, electrons can become self-trapped in the plasma bubble when the wakefield becomes strong enough. This mechanism, known as self-injection, can be further controlled or enhanced using methods such as ionisation injection to achieve localized injection and better beam quality.

An illustration of the bubble regime obtained from numerical simulation is shown in Fig. 1.12.

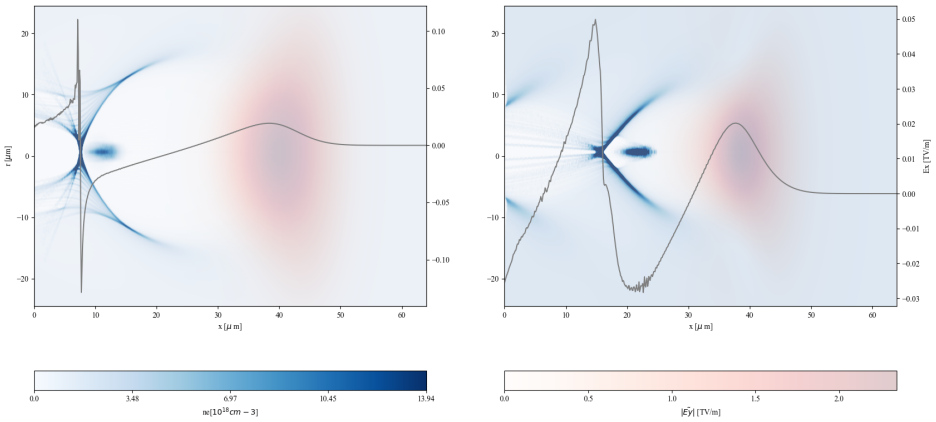

Figure 1.12: Particle-in-cell (PIC) simulation snapshots illustrating the blowout regime of laser wakefield acceleration. The blue colormap represents the electron density (in cm^{-3}), the reddish colormap shows the transverse electric field of the laser pulse (in TV/m), and the gray curve corresponds to the longitudinal electric field (in TV/m). The left panel corresponds to $a_0 = 1.82$ (in vacuum), while the right panel shows a similar configuration with $a_0 = 1.12$ (in vacuum). In both cases, the expulsion of electrons from the laser axis leads to the formation of an ion cavity (bubble). The maximum longitudinal accelerating field reaches 128 GV/m in the left panel and 27 GV/m in the right panel, illustrating the strong dependence of the wakefield amplitude on the laser normalise vector potential a_0.

1.10.1 Injection mechanisms

Numerous injections methods have been investigated in LPA: external injection, colliding pulse injection, ponderomotive injection to name a few cite [11, 12, 13]. In this section we will focus on the three methods that are relevant in our experiment.

Self injection

Self injection [14] occurs in highly non linear wake typically for $a_0 \gtrsim 2$. In the bubble regime some of the electrons forming the sheat of the bubble can be trapped. Indeed if the bubble radius expands fast enough, some of these electrons will be submitted to a transverse focusing force pushing them towards the optical axis as well as longitudinal accelerating gradient. If these electrons gain sufficient "longitudinal" energy during this process (moving at least as fast as the wake) they will be trapped and further accelerated

(see FIG. 1.13).

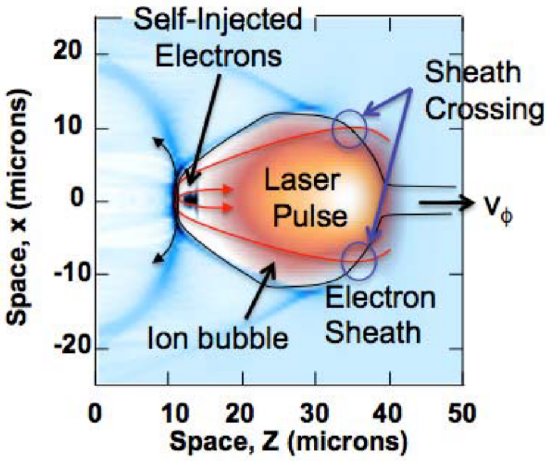

Figure 1.13: PIC code simulation of a laser-wakefield accelerator operating in the blowout or bubble regime. A "matched" 60-fs (FWHM) laser pulse with $a_0 \sim 4$ propagating through a preformed 6×10^{18} cm^{-3} plasma is seen to completely evacuate the region surrounding it of plasma electrons (blue) and leave behind an ion cavity or a bubble (white). Cartoon trajectories of two plasma electrons (black lines) are shown as they stream toward the laser pulse, are then blown out by it, and form part of a negatively charged sheath around the ion cavity. These electrons cross the axis of the laser pulse, overshoot, and form the wake. If the laser pulse is intense enough $a_0 > 4$, some electrons (red lines—cartoon) that started out farther from the laser axis can cross the sheath and gain enough energy from the accelerating field to be injected and trapped inside the potential well of the wake. Adapted from [15].

Ionisation injection

Ionisation injection is the process by which electrons that are ionised at or near the peak of an intense laser pulse are injected in the wake. These electrons can get injected inside the wake because contrary to the background electron they see a non zero potential at their inception, which effectively lower their energy and can allow them to be injected. In order to perform ionisation injection one needs an atom who posses electrons with an ionisation potential close to the peak intensity of the laser pulse. Now let's consider again the condition for trapping set in the above section but this time with an electron born inside the laser pulse. For this electron the transverse momentum and hamiltonian reads:

$$\begin{cases} \mathbf{p}_\perp(\xi) = \mathbf{A}(\xi) - \mathbf{A}(\xi_i) \\ \mathcal{H} = 1 - \phi(\xi_i) \end{cases}$$

With ξ_i the phase at which the electron is ionised. This modification compared to an electron born in front of the laser leads to a new condition for trapping:

$$1 - \phi(\xi_i) \leq \frac{\sqrt{1 + (\mathbf{A}(\xi_{min}) - \mathbf{A}(\xi_i))^2}}{\gamma_p} - \phi_{min} \tag{1.98}$$

By considering the solutions for minimal scalar potential above and considering $\mathbf{A}(\xi_{min}) = 0$ one arrives at the simplified condition in the limit $\gamma_p \gg 1$, $\hat{E}_{max} \gg 1$:

$$\phi(\xi_i) \geq \frac{1}{\hat{E}_{max}^2}$$

To illustrate the condition in terms of a_0 let's consider the linear solution of a squared laser pulse profile [16] with optimal length (inside the pulse):

$$\phi(\xi) = \frac{a_0^2}{8}\left[1 - \cos(\xi) - \frac{\xi}{2}\sin(\xi)\right]$$

This solution leads to $\phi_{min} = \frac{-\pi a_0^2}{8}$ behind the pulse and if the test electron is ionised at the peak of the laser pulse $\phi(\xi_i = -\pi) = \frac{a_0^2}{4}$. We can now express the condition in terms of a_0:

$$1 - \frac{1}{\gamma_p} \leq \phi(\xi_i) - \phi_{min} \leq (\pi/8 + 1/4)a_0^2$$

In the above expression $\mathbf{A}(\xi_i) = 0$ since the ionisation takes place at the peak of the laser pulse, ($\mathbf{E}_{laser} \propto \cos(\xi) = \partial_\xi \mathbf{A}$). In typical LPA setup $\gamma_p > 10$, which means that for ionisation injection to take place a_0 must be greater than 1. As stated before, ionisation injection needs to occur close to the peak of the laser, so for $I \sim 10^{18}[W/cm^2]$ nitrogen atom is a natural choice.

Density-downramp injection

The principle behind *density downramp injection* is to reduce the phase velocity of the plasma wake so that fluid electrons, which were previously too slow to be trapped, can catch up with the wave and be further accelerated.

To understand why the phase velocity decreases in a density downramp, consider how the plasma wavelength changes with density. After a density transition from n_1 to n_2 with $n_2 < n_1$, the plasma wavelength increases:

$$\lambda_{p1} < \lambda_{p2},$$

where λ_{p1} and λ_{p2} are the plasma wavelengths before and after the downramp, respectively.

This increase in wavelength corresponds to a reduction in the local phase velocity of the wake. To quantify this, consider a wake extending N_p plasma periods behind the laser pulse. The distance between a given crest and the laser pulse before the density transition is roughly $N_p\lambda_{p1}$, while after the transition it becomes $N_p\lambda_{p2}$.

The difference in distance,

$$N_p(\lambda_{p2} - \lambda_{p1}),$$

represents the increase in separation relative to the laser pulse due to the downramp. Since this change occurs over the transition length L_{tr}, the corresponding change in phase velocity can be approximated as

$$\frac{\delta v_\phi}{c} = \frac{v_{\phi,2} - v_{\phi,1}}{c} \simeq \frac{N_p(\lambda_{p2} - \lambda_{p1})}{L_{\text{tr}}}, \tag{1.99}$$

where $v_{\phi,1}$ and $v_{\phi,2}$ are the phase velocities before and after the downramp, respectively.

In other words, the slower phase velocity allows electrons that were initially moving too slowly relative to the wave to enter the accelerating phase, enabling *local trapping* in the wake.

1.10.2 Self-focusing

Self-focusing is an important nonlinear effect in laser–plasma acceleration, as it can be used to optically guide the laser pulse through the plasma. Such guiding increases the effective acceleration length and can enhance the maximum energy of accelerated particles. In addition, self-focusing can raise the laser's peak intensity by reducing its transverse spot size.

From the dispersion relation of an electromagnetic wave in plasma, the refractive index can be written as

$$\eta \simeq 1 - \frac{\omega_p^2}{2\gamma\,\omega_0^2}\frac{n}{n_0}, \tag{1.100}$$

For the plasma to act as a focusing medium, the refractive index must be larger on axis and decrease with radial distance r, i.e. $\partial\eta/\partial r < 0$. This is satisfied if γ also decreases with r, meaning $\partial\gamma/\partial r < 0$.

Assuming $\gamma \simeq \sqrt{1 + \mathbf{A}^2}$ since the quiver motion is dominant inside the laser pulse, and noting that the laser intensity is typically peaked on axis and decreases radially, this condition is naturally met. However, it is not sufficient by itself: self-focusing occurs only if the nonlinear focusing overcomes diffraction, which happens when $\frac{P}{P_c} > 1$ [17] Beams in Plasmas), where P is the laser power and P_c is the critical power for relativistic

self-focusing, given in practical units by:

$$P_c\,[\mathrm{GW}] \simeq 17.4\frac{\omega_0^2}{\omega_p^2}. \tag{1.101}$$

The evolution of the spot size $r_s(z)$ can be described by (Sprangle et al.):

$$\frac{r_s^2}{w_0^2} = 1 + (1 - P/P_c)\frac{z^2}{z_R^2}, \tag{1.102}$$

where w_0 is the beam waist in vacuum and z_R the Rayleigh length. This relation shows that the beam diffracts for $P < P_c$, focuses for $P > P_c$, and remains matched (guided) for $P = P_c$.

1.10.3 self-compression

If we consider a Fourier-limited laser pulse driving a wakefield, the pulse can undergo further compression when propagating in a medium with a suitable index of refraction. Let us examine the index of refraction (1.100). Since this quantity depends on the laser intensity through the relativistic Lorentz factor, the pulse experiences self-phase modulation. As a result, the leading and trailing edges of the pulse are red- and blue-shifted, respectively, depending on the local plasma density. The newly generated frequencies enable additional pulse compression.

This effect is evident when considering the local group velocity:

$$v_g \simeq c\eta = c\left[1 - \frac{\omega_p^2}{2\gamma\,\omega_0^2}\frac{n}{n_0}\right].$$

In the blowout regime, the plasma density at the back of the pulse is lower than at the front. Consequently, the rear of the pulse travels faster, leading to temporal compression. Moreover, the regions of highest intensity experience the largest increase in group velocity, further enhancing the self-compression of the pulse.

1.10.4 Limits of laser–plasma acceleration

As discussed in the previous sections, laser–plasma acceleration (LPA) can achieve high accelerating gradients, enabling the production of high-energy electron bunches in a compact structure. However, there are inherent limits tied to the physical processes involved. These limits are primarily characterized by three phenomena: *laser diffraction*, *laser energy depletion*, and *dephasing* between the electron bunch and the wakefield.

We will briefly review each of these limits.

- **Laser diffraction:**

  As discussed in the previous chapter, the spot size of a Gaussian laser pulse in vacuum is given by $w(z) = w_0\sqrt{1+\frac{z^2}{z_R^2}}$, where w_0 is the beam waist and z_r is the Rayleigh length. As the laser pulse propagates through the plasma, its spot size will inevitably increase, thereby reducing the available intensity. This reduction in laser intensity diminishes the ability of the laser to drive the wakefield. Without some form of optical guiding, the wakefield can be maintained for only a few Rayleigh lengths.

- **Laser depletion:**

  As the laser propagates through the plasma, it loses energy by transferring it to the wakefield it drives. The *depletion length* refers to the propagation distance over which the laser pulse transfers all of its initial energy to the plasma wakefield. For example, in the 1D nonlinear theory for a square laser pulse with optimal length, the depletion length is given by [18]:

$$L_{pd} = \begin{cases} \frac{\lambda_p^3}{\lambda^2}\frac{2}{a_0^2}, & \text{if } a_0^2 \ll 1 \\ \frac{\lambda_p^3}{\lambda^2}\frac{\sqrt{2}a_0}{\pi}, & \text{if } a_0^2 \gg 1 \end{cases} \tag{1.103}$$

- **Dephasing:**

  As injected electrons gain energy, they gradually slip out of phase with the plasma wave. Eventually, if accelerated for a sufficient distance, they encounter a decelerating electric field. The *dephasing length* L_d characterizes this distance. In the 1D nonlinear limit, it is given by [18]:

$$L_d = \begin{cases} \frac{\lambda_p^3}{\lambda^2}, & \text{if } a_0^2 \ll 1 \\ \frac{\lambda_p^3}{2\lambda^2}\frac{\sqrt{2}a_0}{\pi}, & \text{if } a_0^2 \gg 1 \end{cases} \tag{1.104}$$

  Depending on the acceleration regime, one of these lengths will act as the limiting factor for maximizing the electron beam energy. For example, if the laser pulse is not guided, the diffraction length will typically be the limiting factor. When guiding is employed, however, diffraction is mitigated, and either the dephasing length or the pump depletion length becomes the primary limitation.

1.11 Electron parameters

A laser-plasma accelerator produces an electron beam characterized by several properties, the relative importance of which depends strongly on the intended application

(e.g., medical, X-ray, etc.). In this section, we provide a brief overview of these key beam parameters.

1.11.1 Electron spectrum

The output electron bunch is composed of many particles that do not all have the same energy see FIG. 1.15. The distribution of their energies is described by the electron spectrum

$$f(E) = \frac{dQ}{dE},$$

from which important statistical quantities can be extracted.

The total bunch charge is given by

$$Q = \int f(E)\, dE.$$

The median energy E_{med} is defined by

$$\int_{-\infty}^{E_{\mathrm{med}}} f(E)\, dE = \frac{Q}{2}.$$

It is generally preferred over the mean energy because the electron spectrum obtained from laser–plasma accelerators (LPAs) is often skewed toward high-energy electrons.

A robust measure of the energy spread is the median absolute deviation (MAD), defined as

$$\delta E_{mad} = \frac{\mathrm{median}\big(\,\big|E - E_{\mathrm{med}}\big|\,\big)}{E_{med}},$$

which characterizes the typical deviation of the particle energies from the median value and is less sensitive to outliers than the standard deviation.

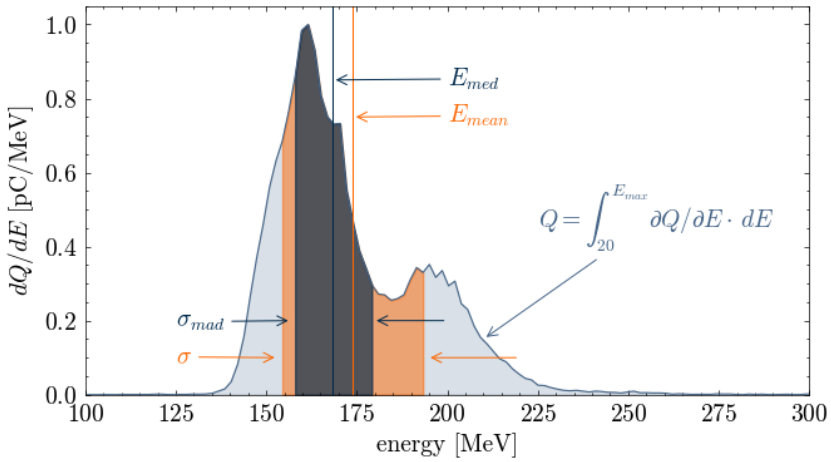

Figure 1.14: electron beam statistical features processed from the electron spectrum

1.11.2 Emittance

The electron beam emittance is a quantity that describes the area occupied by the electron bunch in the projected two-dimensional phase spaces. It is therefore defined in the (x, p_x), (y, p_y), and (z, p_z) spaces and characterizes the spread of the bunch in both position and momentum, making it a general measure of beam quality. An important property of the emittance is that it is conserved under the action of conservative forces, as a consequence of Liouville's theorem, which states that the phase-space density is constant along the trajectories of a Hamiltonian system.

The emittance has several definitions, but the normalized emittance can be expressed as an RMS quantity:

$$\varepsilon_n = \frac{1}{mc}\sqrt{\langle x^2\rangle\langle p_x^2\rangle - \langle x p_x\rangle^2}, \tag{1.105}$$

where $\langle . \rangle$ is the second central moment. This expression can be generalized to the (y, p_y) and (z, p_z) spaces.

1.11.3 Twiss parameters

The classical expression of the emittance is obtained by considering the motion of an electron in a storage ring under the action of conservative forces. Although the magnetic Lorentz force is not strictly conservative, the transverse dynamics can be linearized around the reference trajectory, which leads to an effective Hamiltonian system where the phase-space area is conserved [19].

Assuming that the electron bunch follows a two-dimensional Gaussian phase-space distribution, the motion of an electron located at one standard deviation from the center

of the distribution traces a closed ellipse in transverse phase space. The emittance is then defined as the area of this ellipse and its expression is derived from the periodic motion of the particle in terms of the Twiss parameters (α, β, γ):

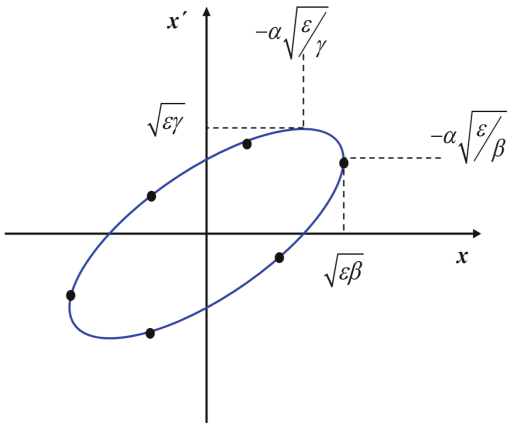

Figure 1.15: The parameters of a phase-space ellipse containing an emittance at a certain point in the lattice. The shape and orientation of the ellipse are determined by the Twiss parameters at the given location. Adapted from [20]

$$\frac{\varepsilon}{\pi} = \gamma x^2 + 2\alpha x x' + \beta x'^2, \tag{1.106}$$

where x is the transverse position and $x' \simeq p_x/p_z$ is the normalized transverse momentum or beam divergence. Each Twiss parameter describes a property of the ellipse:

$$\langle x^2 \rangle = \beta\varepsilon, \quad \langle x'^2 \rangle = \gamma\varepsilon, \quad \langle x x' \rangle = -\alpha\varepsilon.$$

- γ describes how spread the beam is in angle (beam divergence), - β describes how spread the beam is in transverse position, - α measures the correlation between position and angle (tilt of the ellipse).

The two expressions of the emittance can be related through:

$$\frac{\varepsilon}{\pi} = \frac{1}{\bar{p}_z}\varepsilon_n,$$

where $\bar{p}_z$ is the mean longitudinal momentum.

1.11.4 Beam loading

Beam loading arises from the space-charge field of the injected electron bunch. The bunch excites its own wakefield in the plasma, so that the resulting longitudinal electric field is the superposition of the plasma wave driven by the driver and the field generated

by the bunch itself. In the linear regime, and considering the wake of an infinitesimally thin bunch, Katsouleas *et al.* [21] derived the maximum number of electrons that can be loaded while extracting all of the energy from the wave, i.e. completely canceling the accelerating field behind the bunch:

$$N_{\text{max}} = \frac{n_0}{k_p} \frac{E_z}{E_0} A,$$

where $E_z/E_0 < 1$ is the normalized amplitude of the accelerating field and A is the transverse cross-sectional area of the bunch.

At this loading limit, complete extraction of the wake energy comes at the cost of a large energy spread: electrons at the head of the bunch experience the full accelerating field E_z, while those at the tail experience nearly zero field. The energy spread can be reduced by tailoring the bunch profile see FIG. 1.16 . For example, Katsouleas *et al.* showed that a bunch with a triangular longitudinal profile can achieve nearly zero energy spread, albeit at the expense of reduced beam-loading efficiency and a lower accelerating gradient.

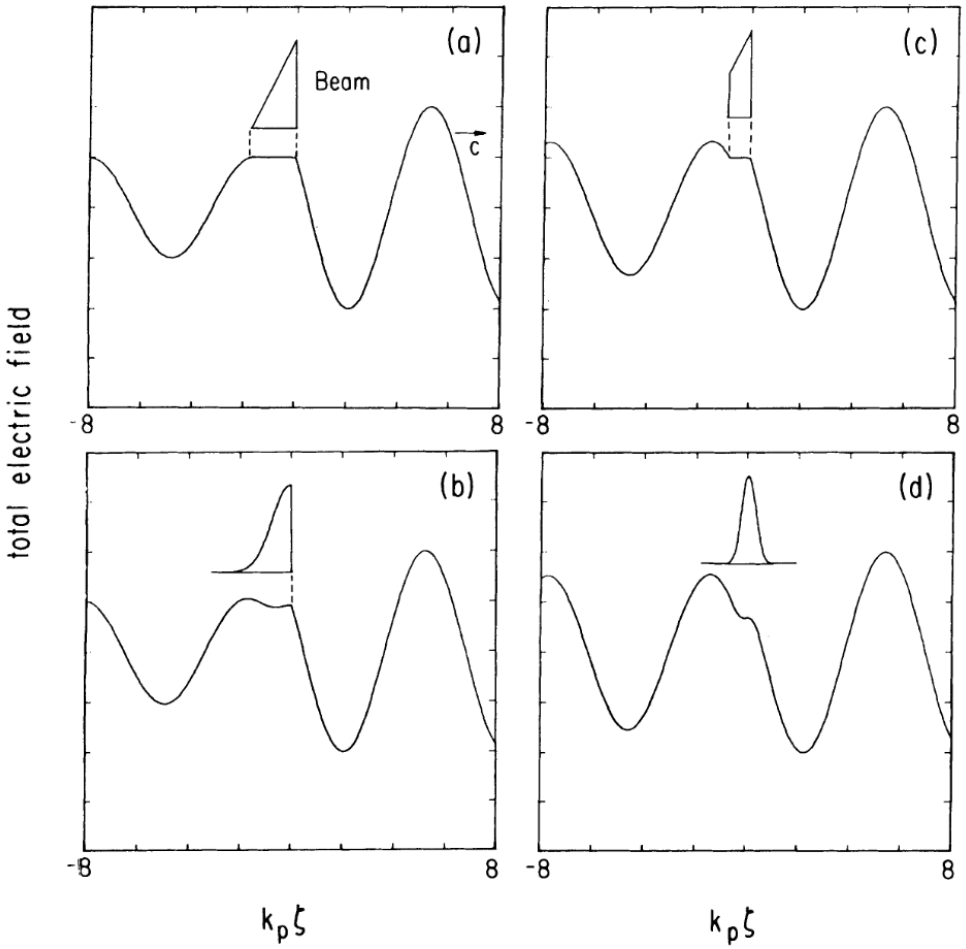

Figure 1.16: Total electric field for various beam shapes: (a) triangle, (b) half-Gaussian of same number of particles, (c) truncated triangle and (d) Gaussian of same number of particle as (c). The shape (a) and (c) leads to the smallest energy spread. Adapted from [21]

References

[1] Donna Strickland and Gerard Mourou. “Compression of amplified chirped optical pulses”. In: *Optics Communications* 56.3 (1985), pp. 219–221.

[2] Peter W. Milonni and Joseph H. Eberly. *Lasers*. 1st. Wiley, 1988, p. 388.

[3] Jean-Claude Diels and Wolfgang Rudolph. *Ultrashort Laser Pulse Phenomena: Fundamentals, Techniques, and Applications on a Femtosecond Time Scale*. 2nd ed. Burlington, MA: Academic Press, 2006. isbn: 978-0-12-215493-5.

[4] Yuto Yamaki et al. “A heuristic approach for actuator layout designs in deformable mirror devices based on current value optimization”. In: *Structural and Multidisciplinary Optimization* 58.3 (2018), pp. 1243–1254. doi: 10.1007/s00158-018-1931-1.

[5] Selcuk Akturk et al. “The General Theory of First-Order Spatio-Temporal Distortions of Gaussian Pulses and Beams”. In: *Optics Express* 13.21 (2005), pp. 8642–8661. doi: 10.1364/OPEX.13.008642.

[6] Spencer W. Jolly, Olivier Gobert, and Fabien Quéré. “Spatio-temporal characterization of ultrashort laser beams: a tutorial”. In: *Journal of Optics* 22.10 (2020), p. 103501. doi: 10.1088/2040-8978/abad08.

[7] Swamp Optics, LLC. *Spatio-Temporal Distortions and How We Measure Them*. https://www.swampoptics.com/spatio-temporal-distortions.html. Accessed: 2026-02-26.

[8] S. V. Bulanov, V. I. Kirsanov, and A. S. Sakharov. “Excitation of ultrarelativistic plasma waves by pulse of electromagnetic radiation”. In: *JETP Letters* 50.4 (1989), pp. 176–178.

[9] V. I. Berezhiani and I. G. Murusidze. “Interaction of highly relativistic short laser pulses with plasmas and nonlinear wake-field generation”. In: *Physica Scripta* 45.1 (1992), pp. 87–94. doi: 10.1088/0031-8949/45/1/015.

[10] W. Lu et al. “Generating multi-GeV electron bunches using single-stage laser wakefield acceleration in a 3D nonlinear regime”. In: *Physical Review Special Topics – Accelerators and Beams* 10.6 (2007), p. 061301. doi: 10.1103/PhysRevSTAB.10.061301.

[11] Barbara Marchetti et al. “Conceptual and technical design aspects of accelerators for external injection in LWFA”. In: *Applied Sciences* 8.5 (2018), p. 757. doi: 10.3390/app8050757.

[12] Jérôme Faure et al. “Controlled injection and acceleration of electrons in plasma wakefields by colliding laser pulses”. In: *Nature* 444.7120 (2006), pp. 737–739. doi: 10.1038/nature05238.

[13] Donald P. Umstadter, Joon-Koo Kim, and E. Dodd. “Laser injection of ultrashort electron pulses into wakefield plasma waves”. In: *Physical Review Letters* 76.12 (1996), pp. 2073–2076. doi: 10.1103/PhysRevLett.76.2073.

[14] S. Kalmykov et al. “Electron self-injection and trapping into an evolving plasma bubble”. In: *Physical Review Letters* 103.13 (2009), p. 135004. doi: 10.1103/PhysRevLett.103.135004.

[15] Chan Joshi. “Laser-Driven Plasma Accelerators Operating in the Self-Guided, Blowout Regime”. In: *IEEE Transactions on Plasma Science* 45.12 (Dec. 2017), pp. 3134–3146. doi: 10.1109/TPS.2017.2769455.

[16] M. Chen et al. “Theory of ionization-induced trapping in laser-plasma accelerators”. In: *Physics of Plasmas* 19.3 (2012).

[17] P. Sprangle, C. M. Tang, and E. Esarey. “Relativistic self-focusing of short-pulse radiation beams in plasmas”. In: *IEEE Transactions on Plasma Science* 15.2 (1987), pp. 145–153. doi: 10.1109/TPS.1987.4316677.

[18] Eric Esarey et al. “Nonlinear Pump Depletion and Electron Dephasing in Laser Wakefield Accelerators”. In: *Advanced Accelerator Concepts Workshop*. Vol. 737. AIP Conference Proceedings. 2004, pp. 578–584. doi: 10.1063/1.1842594.

[19] Helmut Wiedemann. *Particle Accelerator Physics*. 4th. Cham, Switzerland: Springer, 2015.

[20] E. Wilson and B. J. Holzer. *Chapter 2: Beam Dynamics*. Tech. rep. CERN-202. CERN Accelerator School, 2020. url: https://cds.cern.ch/record/2743947/files/Wilson-Holzer2020_Chapter_BeamDynamics_2.pdf.

[21] T. Katsouleas et al. “Beam loading in plasma accelerators”. In: *Particle Accelerators* 22 (1987), pp. 81–99.

2 -PALLAS a Laser-plasma injector test facility

2.1 Introduction

The **PALLAS** project is a laser-plasma injector (LPI) test facility developed within the **EuPRAXIA**[1] initiative, a European program dedicated to advancing plasma-based acceleration technologies. The PALLAS LPI test facility uses the high-energy laser beam from the LASERIX Université Paris Saclay / IJClab facility[2] as a laser driver. The PALLAS project is the first user of the LASERIX facility in terms of beam time and has driven important laser updates during the period of the thesis.

2.2 Objectives

The PALLAS project addresses the critical challenges of controlling and stabilising laser wakefield-accelerated electron beams. The PALLAS LPI project extends beyond the development of an optimised injector for a single application; its primary objective is to establish foundational technological building blocks for future advanced LPI systems. These systems are intended to meet the demanding requirements of applications that require high-quality electron beams, characterised by stability, reliability, and precise control. To achieve these objectives, the project focuses on three key aspects:

- Advanced laser control,
- Plasma target optimisation,
- Electron beam capture and characterisation.

However, basing the project on an existing, relatively simple 40-TW-class chirped-pulse amplification (CPA) laser system imposes inherent limitations, restricting the facility's capacity for exploration and optimisation. To work within these constraints, the PALLAS laser-plasma acceleration (LPA) source uses an ionisation injection technique. This method is chosen because it enables efficient electron injection even with limited laser energy, a critical parameter for producing high-quality electron beams.

2.3 LASERIX - laser driver

The LASERIX Ti:Sapphire laser system serves as the primary driver for the PALLAS laser-plasma interaction (LPI) electron source. This system employs chirped-pulse amplification (CPA) and delivers a peak power of 40 TW at a repetition rate of 10 Hz. Later in this

thesis, a model is presented to predict the behaviour of optical compressor gratings, building upon the work of Leroux et al. Their findings demonstrate that the diffraction grating pair in the pulse compressor deforms during operation at full average power, with comparable power densities. Such deformation adversely affects the laser wavefront, thereby degrading both the spatial and temporal quality of the pulse [3]. Consequently, in alignment with the approach adopted by Leroux et al., the repetition rate was reduced to 1 Hz for the experimental campaign undertaken in this study. The laser system comprises three consecutive amplification stages, which can be categorised into two distinct sections: the frontend and the amplification stages. The system layout, along with the pulse properties following each amplification stage, is depicted in Figure Fig. 2.1.

2.3.1 LASERIX frontend

Pulse seeds generated by a Coherent VITARA mode-locked oscillator are initially amplified using a 10 Hz, four-pass Ti:sapphire (Ti:Sa) BOOSTER amplifier, achieving pulse energies on the order of several microjoules. The amplified pulses are subsequently stretched and directed into a pulse shaper, specifically, a DAZZLER acousto-optic programmable dispersive filter[4]. This device operates by diffracting individual spectral components of the pulse through the excitation of an acoustic wave within a crystal[5]. Such a mechanism enables pre-compensation of phase distortions incurred during subsequent amplification stages, while also facilitating online fine-tuning of the on-target laser temporal pulse characteristics.

At the stretcher output, the laser pulse undergoes further amplification in a regenerative amplifier (REGEN), reaching an energy of 500 µJ. Concurrently, the repetition rate is reduced to 10 Hz via a pulse picker. Both the REGEN and BOOSTER amplifiers are optically pumped using LUMIBIRD CFR and ULTRA flashlamp-pumped, Q-switched lasers.

2.3.2 LASERIX amplification stages

The 10 Hz amplification section comprises three multi-pass amplification stages. Initially, the laser pulse energy is increased to 50 mJ in the first amplifier. Subsequently, the second amplification stage (AMP2), pumped by two flashlamp-based Nd:YAG lasers (PROPULSE1), each delivering 700 mJ, further elevates the pulse energy to 600 mJ. The final amplification stage (AMP3), driven by four flashlamp-based PROPULSE2 lasers, boosts the pulse energy to a maximum of 3 J. To ensure uniform pump beam profiles, the pump laser beams for both AMP2 and AMP3 are spatially homogenised using microlens array homogenisers. Additionally, an attenuator system, comprising a motorised waveplate and two reflective polarisers, enables the adjustment of the laser pulse energy by varying the rotation angle of the half-waveplate. The laser is operated at 1 Hz using the

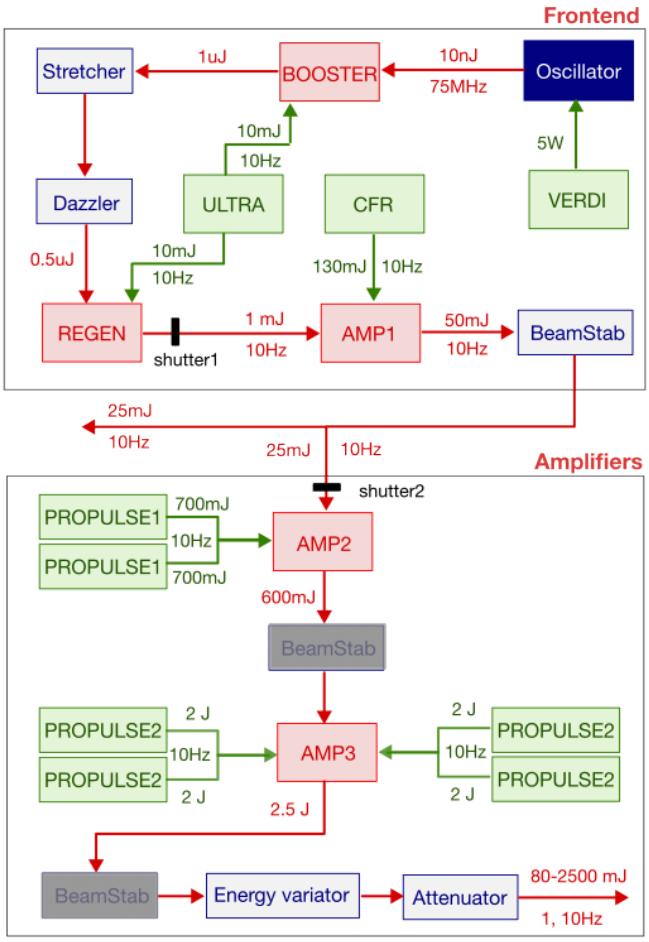

Figure 2.1: Scheme of the LASERIX laser driver system with Frontend module and high energy amplifier module

shutter (shutter2, see Fig. 2.1) . The beam-drift pointing stabilisation due to thermal effects can only operate on the frontend. For the results presented in this doctoral work the beam stabilisation in AMP2 and AMP3 was not used as it requires a fast shutter that must be installed at the output of the laser system and working with large-diameter and high-energy beam.

2.3.3 Laser beam transport line

The laser beam is conveyed from the LASERIX laser facility to the radiation-shielded experimental area via a high-vacuum transport line. Within this beam transport opto-mechanical system, an afocal telescope expands the laser beam diameter from 25 mm to 65 mm. The first lens of this telescope is motorised, permitting precise adjustment of

the longitudinal focal spot position at the target. The total length of the laser beam transport line is ≈ 18.5 m, incorporating turning box chambers, each housing a 45° incidence-angle mirror. The optical transmission efficient of the laser beam transport line is $94 \pm 1\%$.

2.3.4 Optical compressor

The optical compressor is situated within the radiation-shielded experimental area. It employs a folded geometry, incorporating two gold-coated diffraction gratings and a rooftop mirror system (Fig. 2.6). Upon exiting the compressor, the beam is directed towards an in-vacuum deformable mirror (ILAO150) [6]. The optical transmission efficiency of the compressor is approximately $60 \pm 1\,\%$. The two gratings, G1 and G2, are motorised, enabling horizontal-plane rotation. Grating G2 is additionally mounted on a longitudinally motorised stage, which allows adjustment of the output pulse duration from 40 fs to several ps. The laser beam is coupled into the pulse optical compressor with two motorised mirrors. A leak-through of the last mirror is used to monitor near- and far-field.

2.4 LPI accelerator beamline

The laser-plasma injector beamline is structured into three primary modules:

- A **laser injection and focusing module (LIF)**,
- A **plasma module (PLSM)**,
- An **electron characterisation beamline (eCBL)**.

The following sections present the three modules in detail.

2.5 Laser injection module

The laser injection module serves as the central component for controlling the laser driver beam prior to its interaction with the target. Structurally, it integrates a small vacuum chamber equipped with a decoupled vacuum breadboard, all mounted on a massive granite support that also accommodates an optical table dedicated to laser diagnostics. The optical scheme of the LIF module is illustrated in Figure 2.2

Within the vacuum chamber, the module incorporates several key elements: a 1467 mm focal length, 100 mm diameter, $14°$ off-axis parabola (OAP1) for the laser beam focusing; a sampling mirror (LSM) for diagnostic purposes mounted on a motorized stage to direct the beam toward the target; an insertable beam sampler (OBS01) designed for low-energy focal spot measurements ($E < 200$,mJ); and a high-energy laser block (HELB),

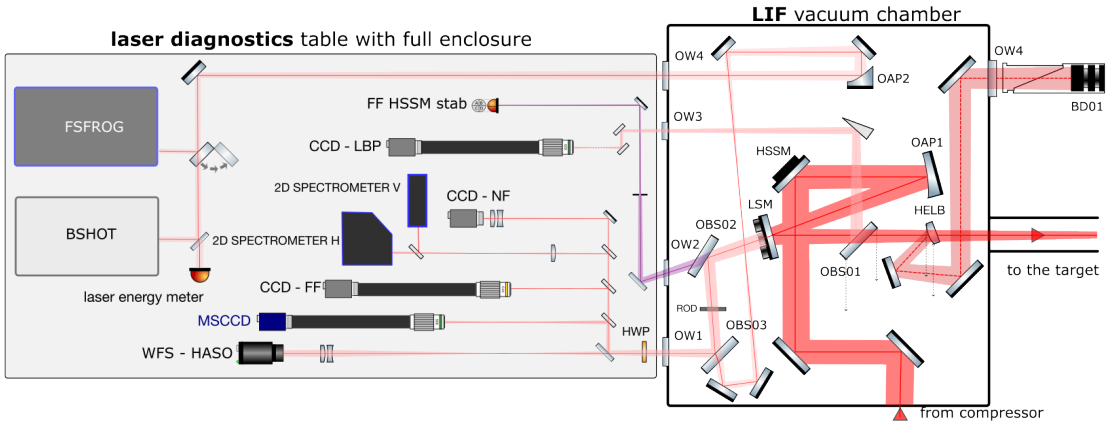

Figure 2.2: Optical scheme of the LIF module

which diverts the beam to a laser beam dump (BD01) during phases of optimisation or characterisation.

Substantial efforts have been dedicated to monitoring laser pulse properties in the vicinity of the interaction region. The Laser injection and focusing (LIF) module incorporates a comprehensive suite of laser diagnostics. The laser sampling mirror (LSM) exhibits an approximate transmission efficiency of 2%. The beam is first spectrally divided at beamsplitter OBS02, where the transmitted component ($\lambda < 780\,\text{nm}$) is directed to a quadrant photodiode for laser beam pointing stabilisation, a system that remained inactive during the course of this doctoral research. Subsequently, the beam is further divided at beamsplitter OBS03, with the transmitted portion re-collimated and directed outside the LIF vacuum chamber for temporal pulse diagnostics, including FROG (Frequency-Resolved Optical Gating)[7], D-SHOT[8], and energy measurement. The beam reflected by OBS03 is then routed to a series of spatio-spectral diagnostic instruments, as detailed in the accompanying table (Tab. 2.1). The deformable mirror, positioned at the compressor output, is imaged onto a wavefront sensor (HASO FIRST) [9]. Operating in a closed-loop configuration, this system enables the deformable mirror to compensate for wavefront aberrations introduced during beam transport, thereby optimising focal spot quality. The sampled focusing beam is directed to a multi-spectral camera (see Chapter 4), mounted on a high-precision motorised linear translation stage, as well as to far-field, near-field, and two cross-imaging spectrometers for comprehensive characterisation

The optical transmission of the LIF up to the target is evaluated to $93 \pm 1\%$, leading to a measured total optical transmission from the output of the LASERIX laser driver system to the target of $52 \pm 1\%$.

Diagnostics	Reference	Resolution	Units
WFS	HASO FIRST	8 ; ~ 2.5	nm; mm
MSCCD	spectral devices msCAM-8B	see section 4	
FF	basler acA640120-gm	0.35	μm
NF	basler acA1920-40gm	$\tilde{7}$0	μm
2D spectrometer H,V	MISS FeamtoEasy	1.5; 6.5	nm; μm
FROG	FS20 FemtoEasy	min 20 (FTL)	fs
D-SHOT	D-SHOT sphereOptics	min 5 (FTL)	fs

Table 2.1: List of diagnostics used in the experiment.

2.6 Plasma module

The plasma module integrates a double-gas-cell target, developed during Pierre Drobniak's thesis [10]. The target is mounted on a motorised hexapod for precise alignment. The gas injection system, comprising four mass flow controllers, enables the use of helium and dinitrogen mixtures with dinitrogen concentrations ranging from 2% to 100%. The system allows independent control of gas pressure in each cell.

2.6.1 Double target cell directly integrated in the beamline

The double-cell target [11] used during commissioning was tested for the first time to generate electron beams during my phd. This original design minimises the compactness of the accelerator elements by placing the two cells directly on the LPA beamline. This arrangement brings constraints on the alignment (no possibility of escaping the target) and differential pumping that must be compacted upward and downward the target (see Fig. 2.6)

The two cells are closed by input and output nozzles that can be customised to adjust the cells' lengths, and in-ramp and out-ramp profiles. For the results presented in this manuscript the cell lengths were both equal to 1 mm and the output nozzle shape was optimised to generate a low-density plateau of a few millimeters at the output.

2.6.2 Gas injection system

The gas flow into the plasma target, as well as the composition of the gas mixture, are regulated in real time using a two-channel mixing device. This system incorporates two Bronkhorst EL-FLOW [12] mass flow controllers (MFC1 and MFC2), each factory-calibrated for specific gas species. These controllers flow the two input gas streams within a turbulent mixer, directing the resultant mixture to a tank reservoir. The achievable mixing ratio depends on the specified output flow rate and is constrained by the operational range of the mass flow controllers. MFC1, calibrated for helium, the primary gas species,

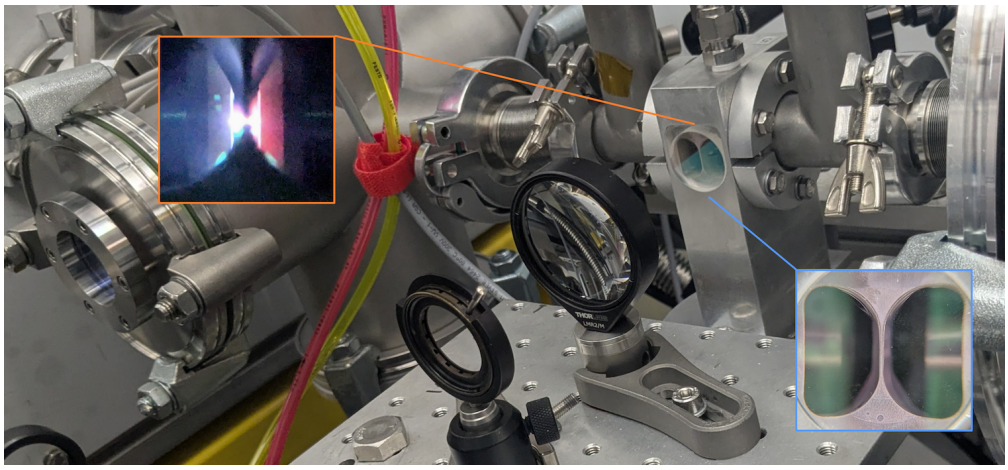

Figure 2.3: View of the double cells target in the beamline with a left-insert zoom on plasma emission with N_2 mixture on the first cell and Helium on the second one. A right zoom insert shows the 2 nozzles with the central wall.

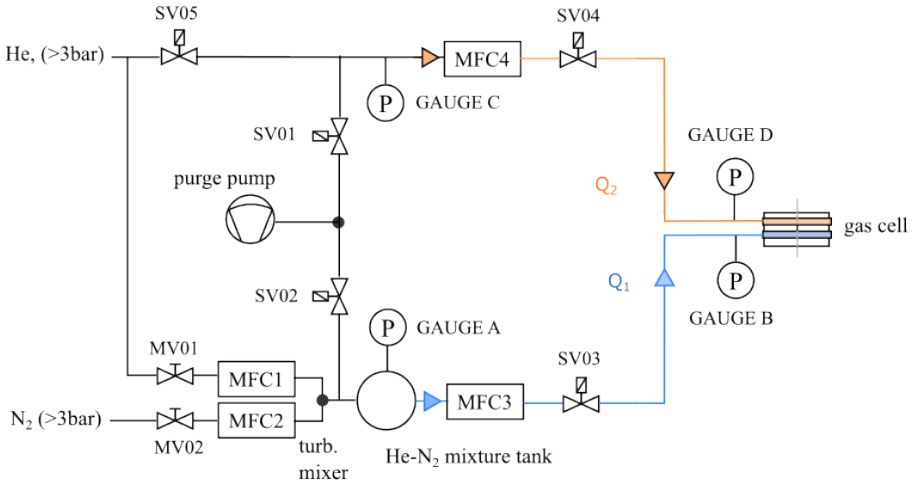

Figure 2.4: Schematic of the gas injection system: MV01, MV02 are manual valves. SV0X are solenoid valves. MFCX are the mass flow controllers. Gauges A and C are pressure transducers (0.1-10 bar range), and gauges B and D are capacitance gauges.

operates over a full range of 1.5 L/min. MFC2, designated for dopant gases, is calibrated for nitrogen and argon and spans a full range of 1.0 L/min. The gas flow injected into the target is further regulated by two additional Bronkhorst EL-FLOW mass flow controllers (MFC3 and MFC4), which manage the gas mixture and pure helium, respectively. Gas pressure is monitored using two types of gauges: capacitance pressure gauges (GAUGES B,D), positioned as close as possible to the target cells for both the gas mixture and pure helium, and pressure transmitters (GAUGES A, C), which measure tank pressure and input pressure within the 1–5 bar range A typical performance of the target profiles is

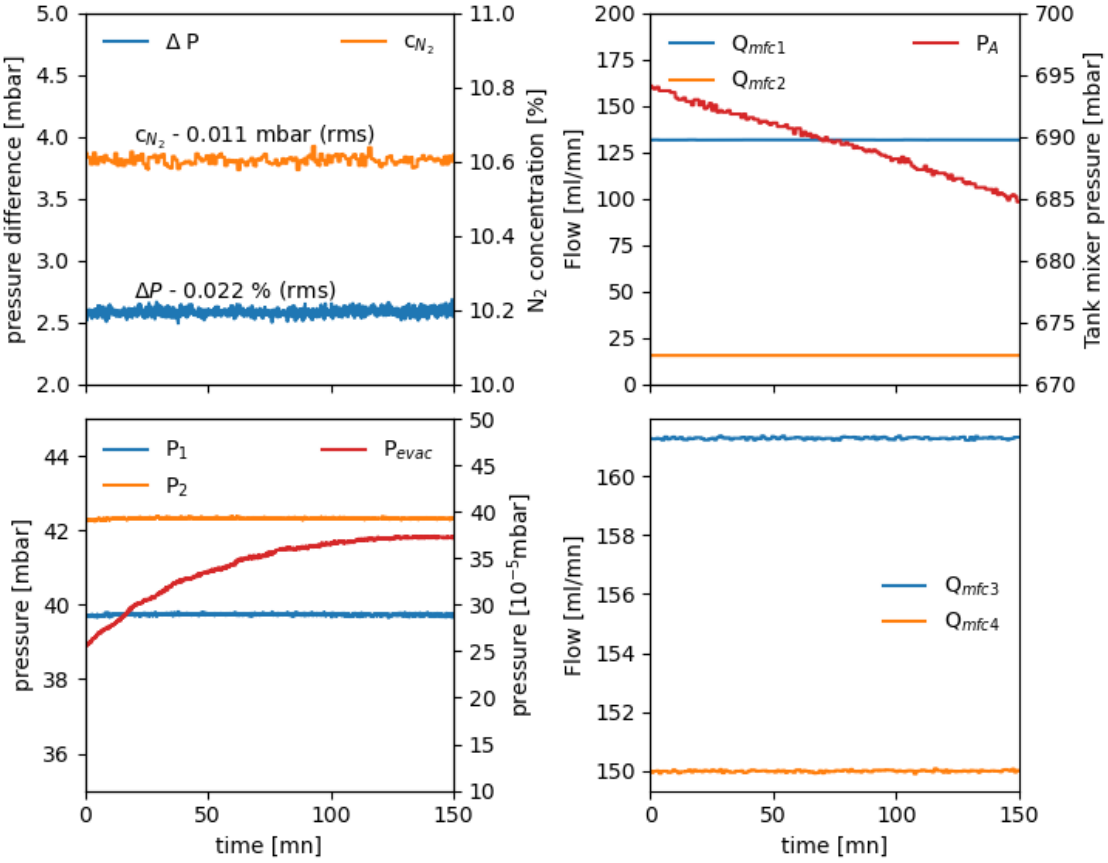

Figure 2.5: Recorded gas injection system parameters showing the stability and vacuum level for more than two hours of operation. The target parameters are set to doping rate (concentration of N_2 of $c_{N_2} = 10.6\,\%$, a pressure of 42 mbar and pressure difference between the two cells of 2.5 mbar

illustrated in the figure (Fig. 2.5) for a fresh target. The pressures p_1 and p_2 are measured by the gauges B and D. The pressure of the mixture Tank is regulated to 0.7 bar with a small drift during the 150 minutes of operation. The mean pressure in the "*evacuation*" vacuum chamber, upward and downward, stays below 4×10^{-4} mbar. These

performances are ideal when the target nozzles are new. After a few days of operation, the entrance and output nozzle apertures are slowly ablated to reach a diameter of 1 mm. Consequently, the flow can be increased by a factor of three to compensate for the increased conductance of the nozzle aperture.

2.6.3 Plasma diagnostics

The plasma target is equipped with three primary diagnostic systems to monitor plasma density, plasma species composition, and target alignment. A SID4-HR wavefront sensor, used in conjunction with a synchronised probe 50 fs femtosecond laser beam, generates two-dimensional density maps of each target cell. The probe delay can be adjusted within a range of -80 ps to $+220$ ps relative to the drive laser pulse. Plasma emission is collected using a high-resolution spectrometer coupled with an intensified gated camera, enabling precise analysis of the dopant localisation between the two cells of the target. Additionally, side imaging, employing a 560 nm bandpass filte, captures visual data of the target. An insertable motorised mirror collects the transmitted drive laser at low energy and relays the central target region onto a camera equipped with a $10\times$ microscope objective. This setup facilitates measurement of target transmission and enables automated alignment.

The plasma target properties are gathered in the following table :

Parameters	**Value**	**Unit**
pressure	5-100	mbar
doping rate (cN2)	2-100	$\%$
pressure difference min; max	-10 ; 10	mbar
plasma electron density cell 1 (*)	$0.5 - 10 \times 10^{18}$	cm^{-3}
plasma electron density cell 2	$0.25 - 4.85 \times 10^{18}$	cm^{-3}
cell length	0.8-2	mm
x,y,z alignement range; accuracy	-5; 2	mm; um

2.7 Electron characterization beamline

The electron characterisation beamline is configured as a free-drift setup without a capture section. The electron spectrometer is positioned 2350 mm from the plasma target.

2.7.1 Laser removal

The laser removal chamber features:

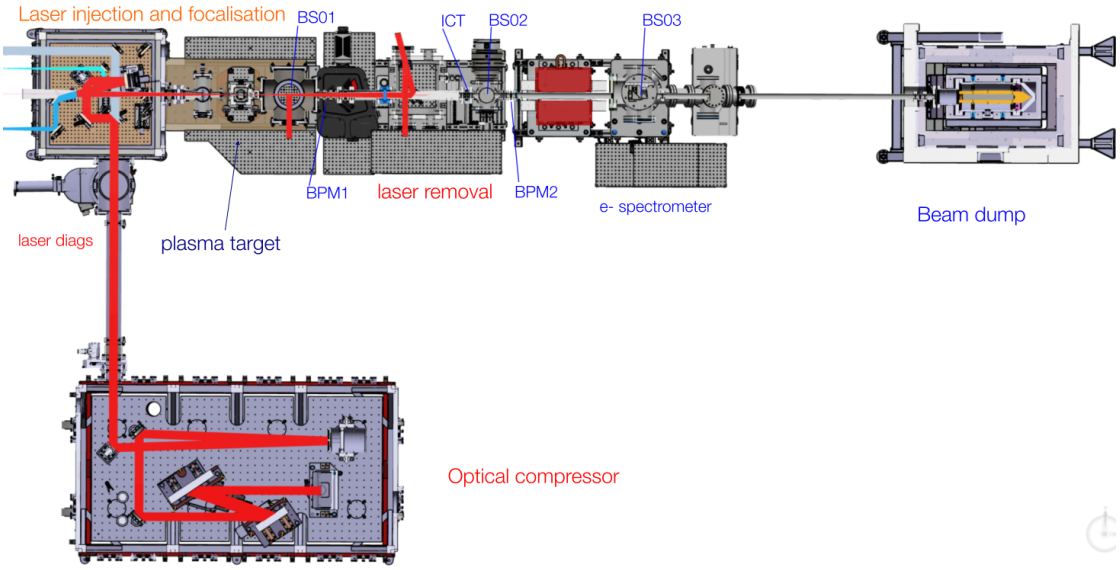

Figure 2.6: Top view of the 3D CAD design of the PALLAS LPA beamline in the radiation shielded area.

- A mirror with a 15 mm diameter central hole to eject most of the laser,
- A 50 µm-thick aluminium foil to block the central part of the laser driver beam.

A small fraction of the laser beam is reflected out of the laser removal chamber onto a diagnostic table that has not been used during this thesis work except for alignment purposes.

2.7.2 Electron beam diagnostics

Beam screens

Three electron screen stations (BS01, BS02, BS03), see FIG. 2.6, are installed to measure the transverse beam profile and pointing at 425 mm and 1790 mm downstream of the plasma target. The BS01 station consists of Lanex screen mounted on the backside of 0.5 mm-thick Aluminium plate positioned at $\sim 70^{\circ}$ incidence of the laser driver/electron axis. The BS01 screen is imaged onto a Basler aca640-120gm 8-bit CCD equipped with a coloured bandpass filter $360 - 580$ nm (FBG39). The BS01 screen is installed on a motorised translation stage and can be removed from beam axis. The BS03 screen station is similar to the BS01 except that the LANEX screen is placed perpendicular to the electron beamline axis and phosphorescence light is collected thanks to a mirror placed

behind the LANEX screen. The BS01 and BS03 screens are mainly used for starting the accelerator and alignment purposes.

The electron beam screen station, designated BS02, was developed by IJClab based on an original design by the Paul Scherrer Institute (PSI) [13]. The primary function of the BS02 station is to acquire beam profile images, thereby facilitating the measurement of beam divergence and enabling precise alignment within the beamline. At its core, the station incorporates a scintillation crystal (YAG:Ce), inserted into the electron beam path via a motorised linear stage. When traversed by an electron beam, the YAG:Ce screen emits visible light at a wavelength of approximately 520 nm, with the signal intensity exhibiting a linear relationship to the beam's charge density [14]. All the Yag:Ce screens have a side Chrome deposited calibration ruler to ease magnification calibration and focusing of the imaging systems.

Table 2.2: Specifications of the BS02 beam screen station.

Parameter	**Value**	**Description**
Screen #1	22 mm diameter $\times$ 1 mm, 18°	YAG:Cer
Screen #2	22 mm diameter	OTR screen (not installed)
Screen #3	12 mm diameter $\times$ 0.1 mm, 0°	YAG:Ce, rear Al coated
CCD camera	2048 $\times$ 2048 px	Basler CMOS acA2040-25gm
Imaging objective	$\times$0.43, WD = 261 mm	MF8M-043-261 macro objective
Spatial resolution	21.1 $\times$ 17.6 µm	
Charge detection limit	$\sim$8 pC	(with Screen #1)

Optically, the screen is imaged from the front at an angle of 18° using two folding mirrors that direct light through a 90° viewport integrated into the vacuum chamber. The scintillation light is subsequently collected via two tilted mirrors and focused through a macro lens objective ($\times 0.43$ magnification) onto a Basler acA2040-25gm CCD camera. The optical configuration adheres to the Scheimpflug principle, ensuring that the entire screen plane is sharply imaged. The imaging system operates with a magnification factor of 0.4. A detailed summary of the individual components and their respective specifications is provided in Table Tab.2.2.

The BS02 station can accommodate up to three distinct screens: two YAG:Ce scintillator screens and one screen designed for optical transition radiation (OTR) measurements. During the current experimental campaign, however, only two YAG:Ce screens 0.1 mm and 1 mm, respectively, have been utilised.

Beam screen and image processing The step-by-step processing of a beam screen image is illustrated in Figure Fig. 2.7 Initially, the raw image of the transverse electron beam profile is corrected by subtracting a background image, captured without an electron signal but including laser light, to eliminate ambient light contamination. The corrected image is subsequently cropped to isolate the YAG:Ce region. Following this, a

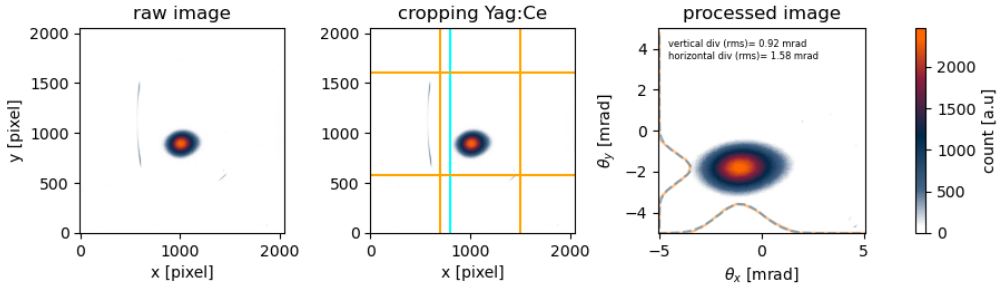

Figure 2.7: Step-by-step processing of a beam screen image.

median filter, averaging the signal over three adjacent pixels, is employed to suppress any remaining hot pixels. The baseline level is extracted from a region of the raw image without an electron signal and subtracted to ensure effective noise removal. The resulting image is presented in figure Fig. 2.7. The final processed image is obtained by cropping to the region of interest and calibrating the spatial dimensions from millimetres to milliradians. To quantify the beam divergence, a two-dimensional Gaussian function is fitted to the processed image, allowing the root-mean-square (RMS) beam divergence to be determined from the projected signal of the fit.

Spatial resolution and charge detection limit The YAG:Ce screen has a total diameter of 25.4 mm, an effective useful diameter of 22 mm, and is available in thicknesses of $100\,\mu$m, $500\,\mu$m, or $1000\,\mu$m. The acceptance angle for freely drifting electrons, from the source to the screen station, is $x = \pm 5.2$ mrad in the horizontal plane and $y = \pm 6.1$ mrad in the vertical plane, based on a distance of 1790 mm from the plasma target. The CCD camera features a chip size of 2048×20482 pixels, yielding a spatial resolution of $21.2 \times 17.6\,\mu$m/pixel. This resolution is influenced by the imperfect alignment of the Scheimpflug angle. While the optical setup provides high spatial resolution, this comes at the cost of reduced charge sensitivity. Depending on the thickness of the YAG:Ce screen, the signal-to-noise ratio per camera pixel is relatively low, resulting in signal loss within the noise. Consequently, only the core region of the transverse intensity profile of the electron beam, where the charge density exceeds the noise threshold, is detectable by the screen station. Figure Fig. 2.8 illustrates the summed pixel counts of the camera image as a function of the total bunch charge, as measured by the integrating current transformer (ICT; see subsection below). A linear fit to the data reveals a zero crossing of the summed image signal at approximately 3 pC, with no detectable charge below 5 pC, thereby indicating the lowest detectable bunch charge. In an ideal diagnostic system without sensitivity limitations, the linear fit would intersect zero at zero bunch charge. However, the data exhibit significant noise and variability, primarily due to beam attenuation caused by the aluminium edge of the first collecting mirror or slight

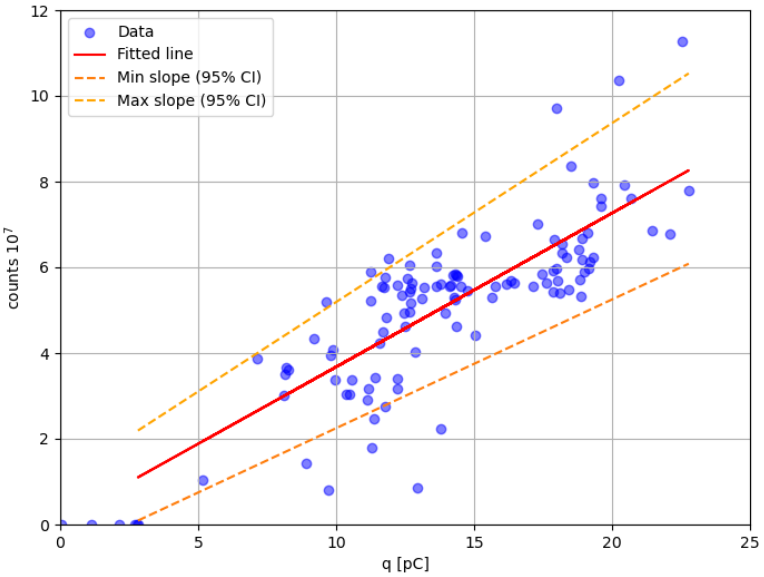

Figure 2.8: Calibration curve for the BS02 beam screen station.

misalignment of the electron beam relative to the screen. Notably, the data presented have not been filtered to exclude clipped beams. Further LPI operations will provide additional data to improve calibration of the BS02 screens. To validate the beam divergence measurements, cross-checking was performed using divergence data obtained from the electron spectrometer.

Integrating current transformer (ICT)

In radiofrequency (RF) particle accelerators, the charge of particle bunches is typically measured non-destructively using current transformers. These devices operate by inductively detecting the beam current, and subsequent integration of the signal yields a value directly proportional to the bunch charge. The proportionality factors, which depend on the specific characteristics of the transformer, can be precisely determined through laboratory calibration or pre-calibration procedures. To ensure the generation of integrable pulses, the current transformer must cover a frequency spectrum that includes near-direct current (DC) components, typically spanning from the kilohertz to the megahertz range. The Turbo-ICT sensor, combined with its associated BCM-RF electronics, enables the accurate measurement of charges for both ultra-short particle bunches and the average currents of continuous-wave (CW) beams. The Turbo-ICT integrates a high-frequency current transformer with a narrow band-pass filter. As de-

scribed in [15], the spectral response, $Q_{out}(f)$ of the current transformer to an incoming current pulse, $I_{in}(t)$, is derived from the product of the pulse's spectral content, $Q_{in}(f)$, and the transformer's transmission coefficient, $S_{21}(f)$. This relationship is expressed as $Q_{out}(f) = Q_{out}(f) \cdot S_{21}(f)$. The time-domain output current pulse, $I_{out}(t)$, is subsequently obtained through the inverse Fourier transform of $Q_{out}(f)$. The amplitude of the output resonance peak is proportional to the input charge and is measured by the BCM-RF electronics, which comprise a narrow band-pass filter and an amplifier. The band-pass filter effectively attenuates noise introduced by the cable connecting the Turbo-ICT to the BCM-RF. When the duration of the input pulse is significantly shorter than the resonance period,a condition satisfied in laser-plasma accelerators (LPAs),the resonance peak becomes independent of other pulse properties. A comprehensive description of the operational principles and calibration procedures for the Turbo-ICT and BCM-RF systems is provided in [15]. Turbo-ICT systems are now widely employed for charge measurements in LPAs[16, 17]. On the PALLAS beamline, this technology enables the detection of bunches with charges as low as 0.1 pC, with an acceptance angle of ± 11 mrad.

Electron spectrometer

The electron spectrometer features:

- A sector dipole electromagnet (max field: 0.8 T),
- Two YAG:Ce scintillating screens (thicknesses: 0.2 mm and 0.5 mm),
- A theoretical energy resolution limit of $\delta E < 0.15$ MeV.

Electron spectrometer design The experimental configuration of the spectrometer deployed at the PALLAS facility, as depicted in Fig. 2.9 is defined by the following technical specifications. At its core, the spectrometer incorporates a sector dipole electromagnet capable of generating a maximum magnetic field of $B_{\text{max}} = 0.8$ T. This configuration produces a curvature radius of 1.15 m and a deflection angle of 17°, corresponding to an effective magnetic path length of 34 cm. The deflection occurs in the vertical plane (y-axis), causing the beam to deflect downward.

Along the deflected beamline, two distinct screen stations are employed for imaging purposes. The first station is positioned 344 mm from the dipole exit, while the second is located 871 mm downstream. Both stations use YAG:Ce scintillating screens, with thicknesses of 0.5 mm and 0.2 mm, respectively (see Fig. 2.9). Each screen measures 80 mm in length and 30 mm in width, and the imaging system is equipped with a camera featuring a pixel resolution of $5.6 \times 5.6\,\mu\text{m}^2$. The dipole magnet itself is situated 2350 mm downstream from the plasma cell exit. A comprehensive summary of the spectrometer's principal parameters is provided in Table 2.3.

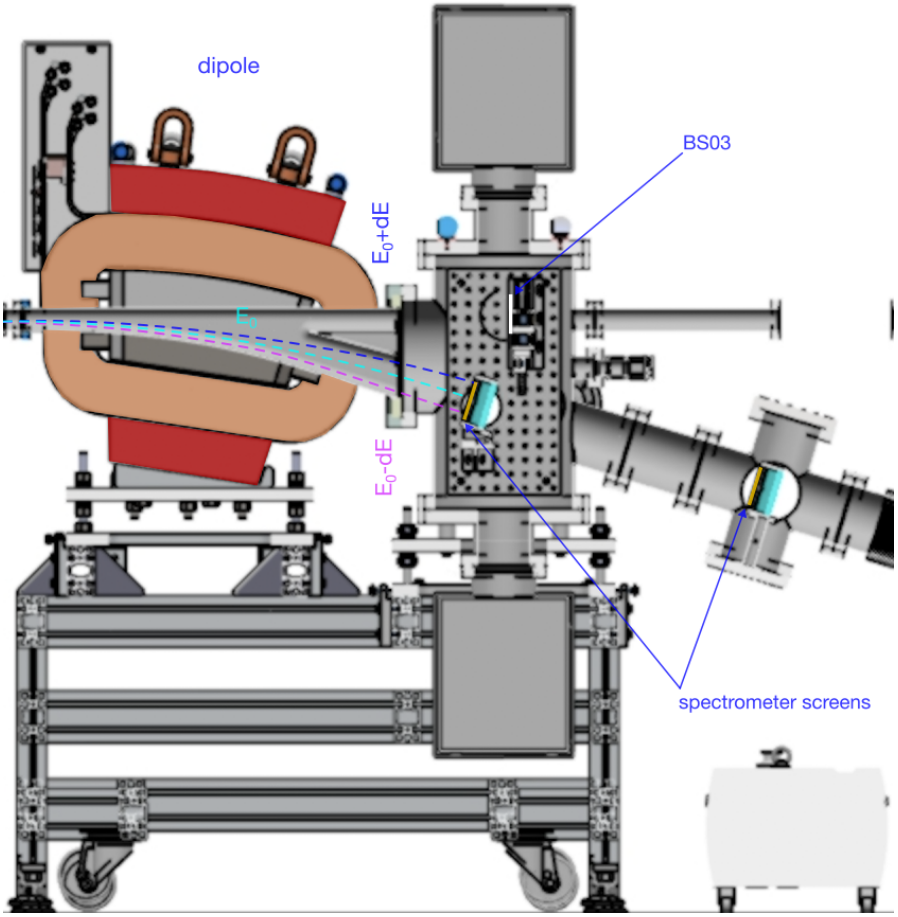

Figure 2.9: Vertical cut of the 3D CAD of the electron spectrometer. The central electron beam trajectory for energy E_0 is represented with the extreme ones $E_0 - dE$ and $E_0 + dE$.

Table 2.3: Specifications of the electron spectrometer

Parameters	**Value**	**Unit**
full range	$20 - 350$	MeV
central energy	$2.374 \cdot I$	MeV
single shot range, $(2 \cdot dE)$	55	$\%$
resolution(*)	~ 0.033; (1.6)	$\%$ (MeV)

Electron spectrometer range and resolution In the final configuration of the PALLAS beamline, the electron beam will be captured and focused using either four quadrupoles or an active plasma lens in combination with two quadrupoles. The focusing strength of these quadrupoles will be adjustable, ensuring that the focal plane will be positioned optimally near the entrance of the dipole magnet.

Additionally, the system will include a collimator assembly located upstream of the dipole. This collimator may be selectively engaged to implement beam energy selection, depending on experimental requirements. Specifically, its deployment will be determined by whether the objective is to analyse the initial beam spectrum or an energy-selected beam spectrum.[18]. The electron spectrometer was designed during Coline Guyot's doctoral research to facilitate the characterisation of electron beams with low energy spread, achieving a theoretical resolution limit of $\delta E < 0.11$ MeV. The energy range is $\pm 27.7\,\%$ of the central energy.
In the first commissioning of the PALLAS beamline performed during my thesis, the electron beamline lacked quadrupole optics, and the beam freely drifted from the plasma target to the spectrometer, limiting resolution to the electron beam's divergence of the electron beam. The central energy of the spectrometer detector„ E_0, is determined by the magnet current according to: :

$$E_0[\text{MeV}] = 2.37 \cdot I[\text{A}]$$

For the range of energy of interest for PALLAS, $20 - 350$ MeV, the formula of energy distribution on the first spectrometer screen, with $x = 0$ set at the center of the screen, is

$$E(x, I) = (-0.0000417 \cdot x^2 + 0.017 \cdot x + 2.3740) \cdot I \approx (2.3740 + 0.017 \cdot x) \cdot I$$

where $E(x, I)$ is in MeV, x in mm and I in A.

The first electron spectra obtained are presented in the commissioning results section. It is important to note that the electron spectrometer was specifically designed for operation in conjunction with a capture-and-focusing system, such as a quadrupole triplet. Consequently, its deployment in free-drift mode is less optimal compared to a standard configuration, which typically employs a permanent dipole magnet positioned immediately downstream of the plasma target and coupled with a scintillator screen for detection.

Electron spectrometer image processing The scintillation light emitted by the spectrometer screen is collected using two Basler acA4096-11gm CCD cameras, each featuring a chip size of 4112×2176 pixels. These cameras are equipped with macro lens objectives (MF8M-0247-267; Lensation). The resolution is $14.2\,\mu$m/pixel. The two cameras—designated as *spec1* and *spec2*—are positioned such that each covers slightly more than half of the screen. The folded mirror positioned behind the scintillation screen is motorised for the first screen. The image processing pipeline, developed by another PhD student within the group, Mykyta Lenivenko, leverages GPU acceleration to meet the speed requirement of less than 50 ms, thereby enabling 10 Hz image acquisition and processing from both CCDs. The workflow comprises the following steps:

- **Preprocessing**: Conducted once on illuminated screens to define cropping coordinates and correction angles (required whenever the system is realigned, e.g., during screen changes). These parameters are saved to a JSON configuration file.
- **Background Correction**: Applied to both images.
- **Geometric Correction and Concatenation**: Both images undergo angle correction, cropping, and concatenation on the GPU.
- **Data Transfer and Storage**: The concatenated image is transferred to the CPU and saved to disk.

Sensitivity studies of both detectors reveal consistent responsivities, though these are not homogeneous along the horizontal direction of the chipset.

2.8 First results from PALLAS electron source commissioning

Over a three-month period, the complete PALLAS experimental setup,from the laser system to the electron characterisation beamline, underwent commissioning. Following one month of laser parameter optimisation, the first electron signals were successfully detected at the initial screen station (BS01). Subsequently, parallel efforts were dedicated to further refining laser beam parameters and commissioning the electron beam characterisation beamline. The data presented herein were acquired at the conclusion of the commissioning phase, prior to the final amplification stage (AMP3) upgrade of the LASERIX laser driver. Instrument control and command were implemented through the TANGO Controls open-source distributed control system[19]. Data acquisition was facilitated using custom archivers developed on top of the HDB++/TimescaleDB TANGO archiving framework[20]. At the time of manuscript preparation, the initial commissioning results remain under active analysis and will be reported in detail in forthcoming publications.

2.8.1 Laser properties

This section presents a summary of the optimised temporal and spatial laser pulse properties that facilitated the generation of the electron beam during the commissioning of the electron characterisation beamline. The experimental data were acquired using the LIF diagnostics setup, as described in the preceding section, with measurements performed using the LBP and FS-FROG systems.

Spatial laser pulse properties

For the electron beam measurement presented here, closed loop operation of the deformable mirror and the wavefront sensor was used to optimise the wavefront and a Strehl ratio of 0.62 ± 0.05 was achieved at high energy at repetition of 1 Hz. The corresponding focal spot measured with LBP diagnostic are presented in the figures (Fig. 2.10 , 2.11) and report a lower Strehl ratio 0.45 ± 0.03.

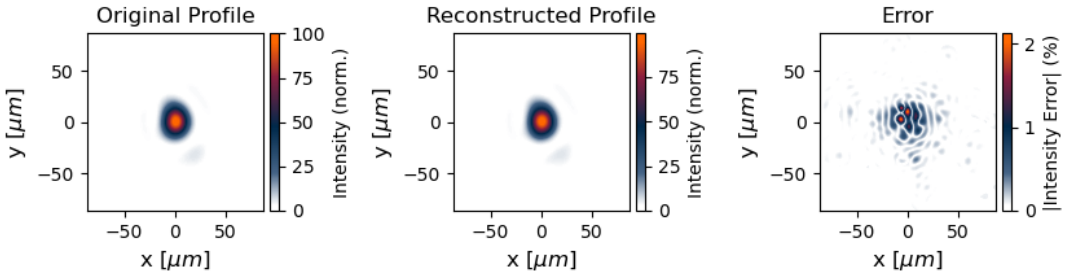

Figure 2.10: High energy focal spot - *(left)* original focal spot image with background substraction recorded with LIF LBP diagnostic - *(center)* - cleaned image using lasy [21] profiles using $N = 20$ HG modes - *(right)* reconstruction error

Using LASY librairy functions[21] the images have been cleaned and denoised (Fig. 2.10 , 2.11). The measured profiles are decomposed into Hermite-Gauss modes. The cleaning is done by keeping only the first modes here, $N = 20$. The focal spot width at $1/e^2$ are 20.7 ± 0.2 um and 21.2 ± 0.1 um along horizontal $x-$axis and vertical $y-$axis respectively. In the figure Fig. 2.11, the intensity profiles range is truncated by the limited dynamic range (typically 40 dB) of the LBP CCD detector encoded on 12 bits bit-depth. The online monitoring of the far field at high energy was not available during the July 2025 due to a not well adjusted sampling level of beamsplitter OBS01/OBS02 in the LIF.

This default has been corrected for the experimental campaign presented in chapters 4 and 5.
The laser beam pointing stability at the output of the amplifier (AMP3) and similar to the one of the output far-field of the compressor (not presented here because less accurate) is shown in figure Fig. 2.12 :

The beam pointing stability is around a radius of the focal spot in RMS but extend to three times the focal spot in the vertical direction mainly due to the thermal drift in the laser system not compensated during this measurement campaign. The beam pointing instability and low Strehl ratio limit the target nozzle lifetime and electron beam stability.

Temporal laser pulse properties

The temporal profile of the laser was measured at low energy with the FROG[7] and B-SHOT[8, 22]. The laser pulse duration (FWHM) is evaluated 42 ± 4 fs and 48 ± 4 fs using the FROG and D-SHOT respectively. The energy level was a bit low for both diagnostics leading to noisy measurements. The spectral phase reconstruction of the FS20 FROG failed to reconstruct the fundamental spectrum properly. In the following, a pulse duration of 45 fs is considered.

Laser energy stability

The output laser energy is recorded for every shots. The laser is operated by burst of 20-30 shots at 1 Hz. Stacked data are presented in the figure (Fig. 2.13). Beam pointing stabilisation was only activated in front end of the laser (see Fig. 2.1). No re-alignment was done in the amplifier AMP2 and AMP3. The laser energy exhibits a stability of $1\,\%$ (RMS). The drop around the shot 2000 is due to a long pause ($\sim$ 45 mn) in the firing sequence impacting the state of the laser system.

In the following table the laser parameters on target are summarised in the table :

parameters	**value**	**unit**
energy, E_l	0.85 ± 0.08	J
peak intensity, I_l	4.63×10^{18}	W/cm^2
normalised potential vector, a_0	1.47	
beam waist radius $w_{0,x}$; $w_{0,y}$	10.10 ± 0.1 ; 11.35 ± 0.1	um
pulse duration (FWHM), τ_l	45 ± 4	fs
Strehl ratio	0.45 ± 0.03 - (0.62 ± 0.05)	

2.8.2 Electron beam properties

The commissioning of the LPA electron source and beamline primarily focused on refining alignment procedures, instrument tuning, and calibration of electron diagnostics,

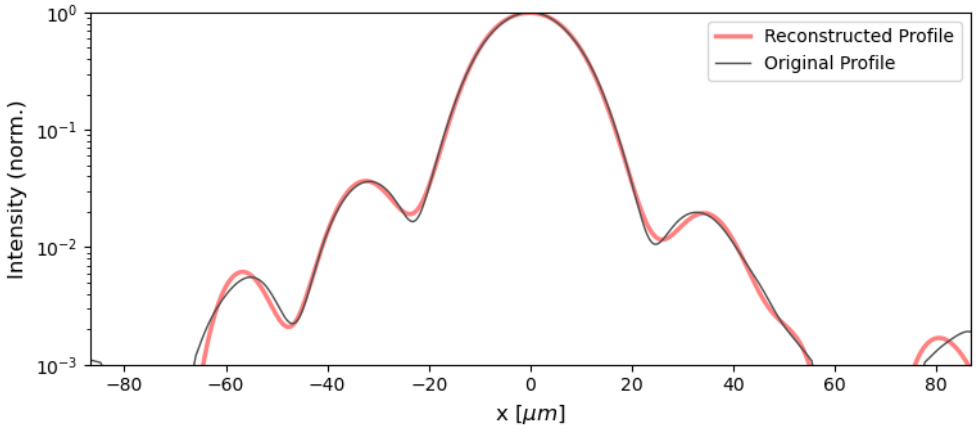

Figure 2.11: High energy focal spot horizontal profiles reconstructed using lasy using $N = 20$ HG modes

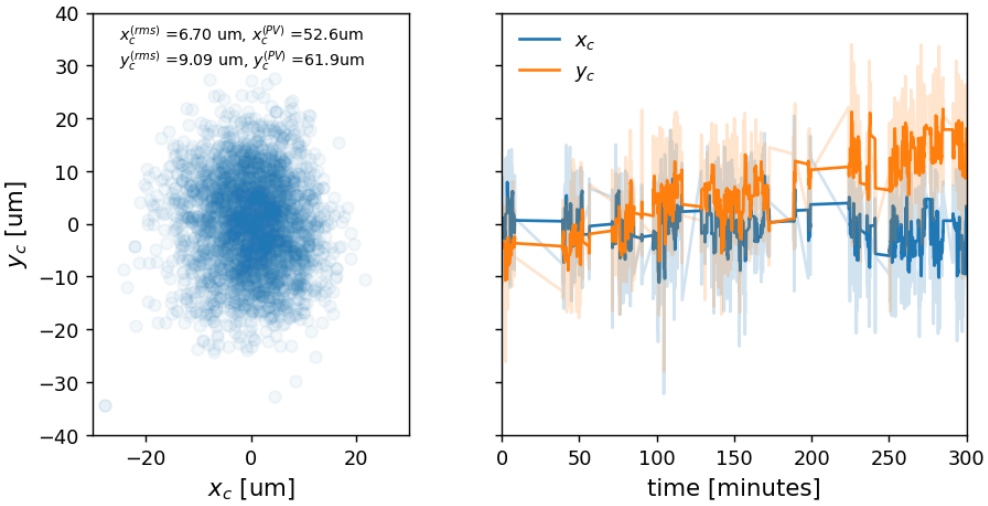

Figure 2.12: Recorded laser pointing stability at the output of the last laser amplifier (AMP3) during last commissioning run for every shots fired to generate electrons. - *(left)* centroid position of the output far field of the amplifier (AMP3) - *(right)* time evolution of the centroid (x_c,y_c)

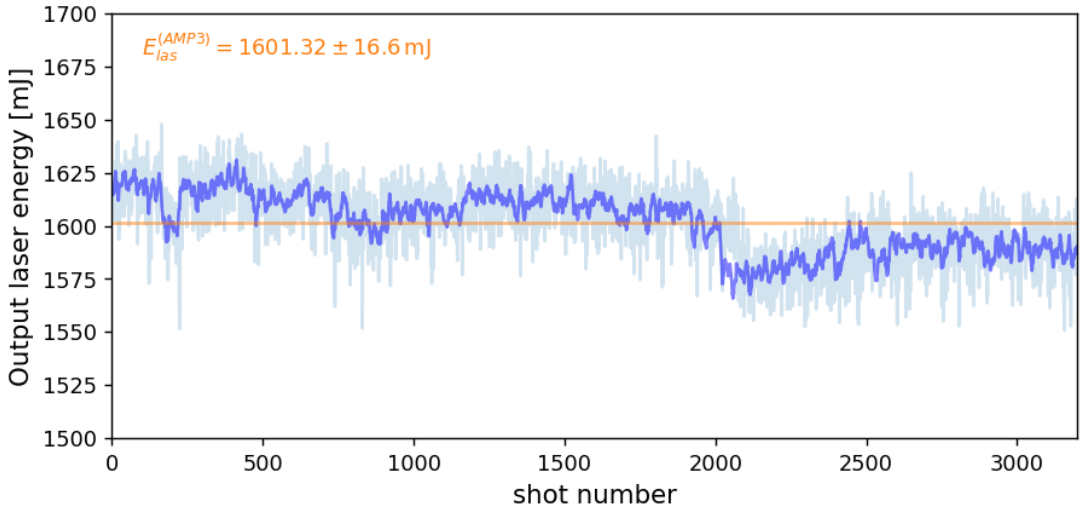

Figure 2.13: Recorded laser energy stability at the output of the last laser amplifier (AMP3) during last commissioning run.

rather than on the direct optimisation of the electron beam itself. As the beamline is not equipped with corrective dipole optics, the alignment of the electron beam is achieved using the final laser optical mirror (LSM; see Fig. 2.2) and the adjustable target. The target can be repositioned to modify the beam angle, ensuring proper propagation to the spectrometer, which is located over 2 meters from the source. To calibrate the electron diagnostics, a baseline optimisation of the electron beam, with respect to pointing stability, divergence, and charge, was necessary prior to its entry into the spectrometer.

The following electron beam properties measured and presented here have been obtained after a quick scan optimisation leading to a stable charge beam around 20 pC and divergence and pointing stability allowing the tuning of the electron spectrometer . The plasma parameters are summarised on the following table : The electron beam

parameters	**value**	**unit**
p_1	73.38 ± 0.87	mbar
$\Delta P = p_2 - p_1$	3.5 ± 0.32	mbar
c_{N_2}	38 ± 0.4	%
plasma electron density cell 1 (*)	8.97×10^{18}	cm^{-3}
plasma electron density cell 2	3.73×10^{18}	cm^{-3}
cell length	1; 1	mm
telescope lens position	-2	mm

Table 2.4

exhibits an asymmetry in the plane of laser polarisation (horizontal axis), which is per-

pendicular to the dispersion axis of the spectrometer. Beam pointing data, recorded over several shot sequences at the BS02 screen station following realignment, demonstrate limited stability but retain an asymmetry consistent with that of the laser (see Fig. 2.14). Despite this, the beam divergence remains constrained, highlighting the critical role and robustness of an optimised plasma electron density profile and the presence of a low-density plateau at the exit of the second target [23]. Once the spectrometer

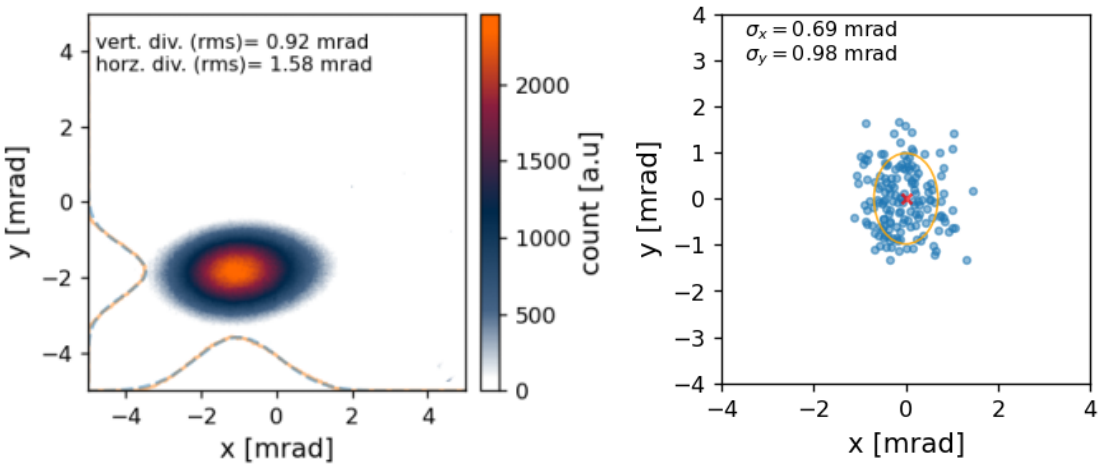

Figure 2.14: *(right)*- typical electron beam profile obtained after aligning the electron beam on BS02. *(right)* - electron beam pointing stability on the electron beam screen BS02.

is aligned, the measurable energy range will extend to $\pm 27.5\%$ around the detector's central energy E_0. A complete current scan was performed to vary the deflection and set the current to the correct energy range. Upon alignment, the measurement range is centred on the electron beam's energy peak. A finer scan, focused around 140 MeV, is presented in Fig. 2.15 with an energy range of approximately 65 MeV in the considered range of current applied to the dipole. The mean energy is 141.8 ± 7.8 MeV with an accuracy limited by the divergence, which gives, considering 1 mrad of pointing jitter and 1 mrad divergence, leading to an error of ~ 1.6 MeV. The energy spread evaluated using the median absolute deviation is 12.05 ± 1.87 MeV.

A broader scan enables reconstruction of the full energy spectrum, ranging from 20 to 220 MeV. Making 10 consecutive shots while scanning for various dipole currents of the spectrometer, covering a central energy range from 20 MeV to 210 MeV, yields the reconstructed spectrum shown in figure (Fig. 2.16).

The spectrum show no low energy components. The stitching remains to be optimised to retrieve the full spectrum. This feature is important when starting optimisation with a broader electron spectrum.

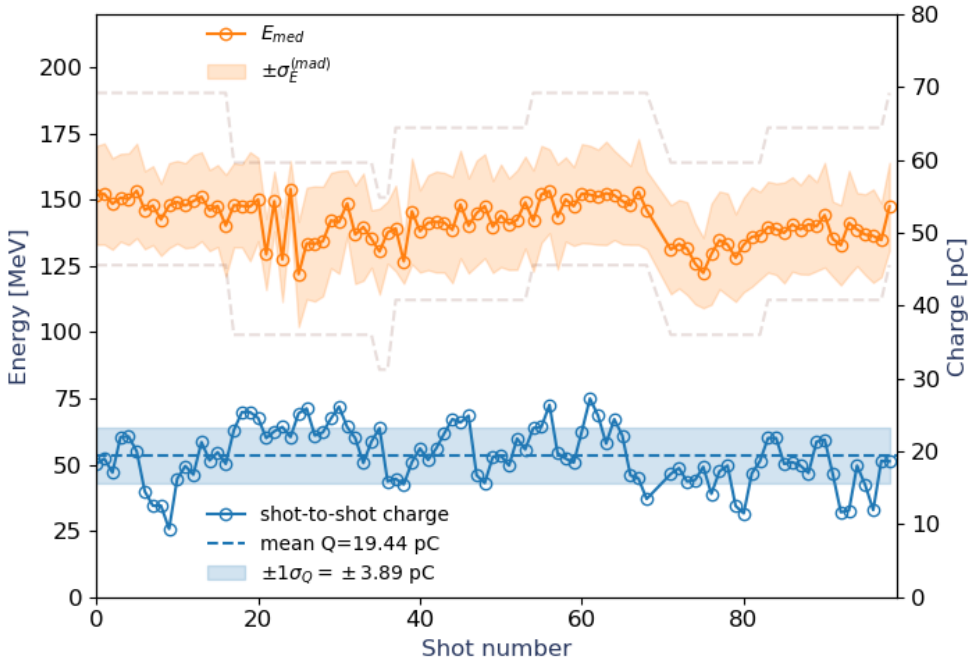

Figure 2.15: Selected sequence obtained with target parameter presented in table Tab. 2.4 showing electron beam energy E_{med}, median absolute deviation $\sigma_E^{(mad)}$ and charge.

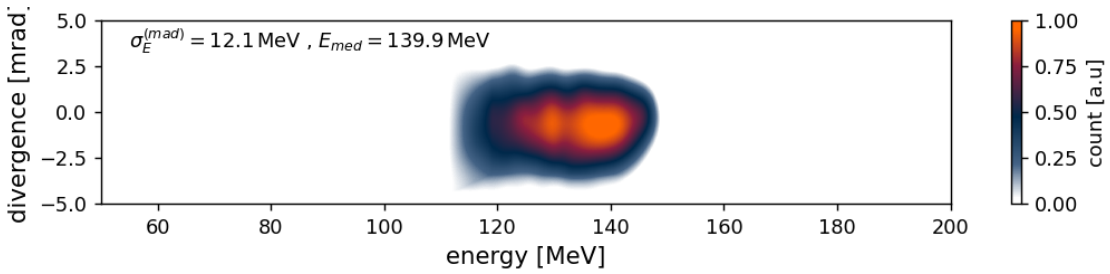

Figure 2.16: An example of a reconstructed spectrum with measurements taken from 20 MeV to 210 MeV.

The key properties of the resulting electron beam are summarised in the following table:

parameters	**value**	**unit**
E_{med}	141.8 ± 7.8	MeV
E_{mad}	12.05 ± 1.87	MeV
q	19.4 ± 3.9	pC
divergence (horz. $\times$ vert.)	$(2.72 \pm 0.48) \times (1.73 \pm 0.95)$	mrad

Despite the limited time allocated to the optimisation of the electron beam, satisfactory parameters were nevertheless achieved during this initial experimental campaign. However, several critical aspects remain to be addressed: these include the refinement of laser diagnostic systems (as discussed in Chapters 4 and 5), the enhancement of focal spot quality, and the improved control of the longitudinal focal position. The latter parameter, in particular, proved to be the most critical factor in successfully generating an electron beam and the onset of acceleration. The absence of an active laser stabilisation system, compounded by the lack of a high-energy shutter at the output of the LASERIX laser, further constrained the pointing stability and energy consistency of both the drive laser and the produced electron beam. Additionally, the sensitivity of the electron spectrometer requires further improvement, as the signal-to-noise ratio (SNR) for the presented data remained relatively low (~ 7). The plasma parameters employed during these experiments corroborate the limited intensity of the laser driver, necessitating operation under conditions of high nitrogen (N_2) doping and elevated gas pressure. Although these conditions deviate from the theoretically optimal parameter ranges, the resulting electron beams nevertheless exhibited acceptable performance characteristics.

References

[1] R. W. Assmann et al. “EuPRAXIA Conceptual Design Report”. In: *The European Physical Journal Special Topics* 229.24 (Dec. 2020), pp. 3675–4284. issn: 1951-6355, 1951-6401. doi: 10.1140/epjst/e2020-000127-8. (Visited on 03/01/2021).

[2] Elsa Baynard et al. “Laserix : Une plate-forme pour le développement large spectre d’applications de la lumière intense”. In: *Photoniques* 118 (2023), pp. 30–34. issn: 1629-4475, 2269-8418. doi: 10.1051/photon/202311830. (Visited on 03/06/2026).

[3] Vincent Leroux et al. “Wavefront Degradation of a 200 TW Laser from Heat-Induced Deformation of in-Vacuum Compressor Gratings”. In: *Optics Express* 14.10 (2018), p. 11.

[4] FASTLITE AMPLITUDE. *DAZZLER - ultrafast pulse shaper*. https://amplitude-laser.com/fr/produits/lasers-femtosecondes/instrumentation-fr/dazzler/.

[5] Pierre Tournois. “Acousto-Optic Programmable Dispersive Filter for Adaptive Compensation of Group Delay Time Dispersion in Laser Systems”. In: *Optics Communications* 140.4 (1997), pp. 245–249. issn: 0030-4018. doi: 10.1016/S0030-4018(97)00153-3.

[6] Imagine Optics. *ILAO STAR 150*. https://www.imagine-optic.com/products/ilao-star-deformable-mirrors/.

[7] FemtoEasy. *Single shot Fast FROG FS20*. https://www.femtoeasy.eu/product/frogs/single-shot-fast-frog/.

[8] Sphere Ultrafast photonic. *Dispersion Scan - BSHOT*. https://www.sphere-photonics.com/b-shot-single-shot-pulse-temporal-measurement/.

[9] Imagine Optics. *Wavefront sensor HASO FIRST*. https://www.imagine-optic.com/products/haso4-first-wavefront-sensor/.

[10] P. Drobniak. “Développement d’un Prototype d’injecteur Laser-Plasma à 150 MeV”. PhD thesis. Orsay, France: Université Paris-Saclay, Dec. 2023.

[11] P. Drobniak et al. *Two-chamber gas target for laser-plasma accelerator electron source*. 2023. arXiv: 2309.11921 [physics.acc-ph]. url: https://arxiv.org/abs/2309.11921.

[12] Bronkhorst. *EL-FLOW select mass flow controller*. https://www.bronkhorst.com/products/gas-flow/el-flow-select/.

[13] Rasmus Ischebeck et al. “Transverse Profile Imager for Ultrabright Electron Beams”. In: *Physical Review Special Topics - Accelerators and Beams* 18.8 (Aug. 2015), p. 082802. issn: 1098-4402. doi: 10.1103/PhysRevSTAB.18.082802. (Visited on 01/30/2024).

[14] Thomas Kurz et al. “Calibration and Cross-Laboratory Implementation of Scintillating Screens for Electron Bunch Charge Determination”. In: *Review of Scientific Instruments* 89.9 (Sept. 2018), p. 093303. issn: 0034-6748, 1089-7623. doi: 10.1063/1.5041755. (Visited on 07/24/2023).

[15] F Stulle et al. “Turbo-ICT Pico-Coulomb Calibration to Percent-level Accuracy”. In: (2015), p. 4.

[16] K Nakamura et al. “Pico-Coulomb Charge Measured at BELLA to Percent-Level Precision Using a Turbo-ICT”. In: *Plasma Physics and Controlled Fusion* 58.3 (Mar. 2016), p. 034010. issn: 0741-3335, 1361-6587. doi: 10.1088/0741-3335/58/3/034010. (Visited on 03/03/2020).

[17] T. André et al. “Control of Laser Plasma Accelerated Electrons for Light Sources”. In: *Nature Communications* 9.1 (Dec. 2018), p. 1334. issn: 2041-1723. doi: 10.1038/s41467-018-03776-x. (Visited on 01/15/2020).

[18] Coline Guyot. “Optimisation of Electron Beam Performance for High Peak Current Laser-Plasma and Multi-Pass Energy Recovery Accelerators with 6D Tracking Start-to-End Simulations”. PhD thesis. Orsay, France: Université Paris-Saclay, 2024.

[19] Scott Feister et al. “Control Systems and Data Management for High-Power Laser Facilities”. In: *High Power Laser Science and Engineering* (June 2023), pp. 1–31. issn: 2095-4719, 2052-3289. doi: 10.1017/hpl.2023.49. (Visited on 07/24/2023).

[20] Lorenzo Pivetta et al. “New Developments for the HDB++ TANGO Archiving System”. In: *Proceedings of the 16th Int. Conf. on Accelerator and Large Experimental Control Systems* ICALEPCS2017 (2018), 5 pages, 2.075 MB. doi: 10.18429/JACOW-ICALEPCS2017-TUPHA166. (Visited on 02/02/2022).

[21] Maxence Thévenet et al. *LASY: LAser Manipulations Made eaSY*. Mar. 2024. doi: 10.48550/arXiv.2403.12191. arXiv: 2403.12191 [physics]. (Visited on 03/05/2026).

[22] Chen Guo et al. “Single-shot, high-repetition rate carrier-envelope-phase detection of ultrashort laser pulses”. In: *Opt. Lett.* 48.20 (Oct. 2023), pp. 5431–5434. doi: 10.1364/OL.498664. url: https://opg.optica.org/ol/abstract.cfm?URI=ol-48-20-5431.

[23] Y.-Y. Chang et al. “Reduction of the Electron-Beam Divergence of Laser Wakefield Accelerators by Integrated Plasma Lenses”. In: *Physical Review Applied* 20.6 (Dec. 2023), p. L061001. issn: 2331-7019. doi: 10.1103/PhysRevApplied.20.L061001. (Visited on 12/11/2023).

3 -Surrogate Models studies for laser-plasma accelerator

3.1 Motivation

This section results are based on an article that I wrote as first author with the addition of further results on spectrum reconstruction using Machine learning technique and a foreword on the construction of reverse machine learning models.

This work continues the efforts initiated by the former PHD student of the group Pierre Drobniak [1, 2]. In his study, a random scan of Particle-In-Cell (PIC) simulations was performed over the laser and plasma parameters available at the PALLAS platform using the SMILEI code [3] to better characterize the accessible parameter space. This approach made it possible to identify target operating points that satisfied the PALLAS requirements, namely, electron beams with energies in the range of 150–250 MeV, charges of 30–50 pC, energy spreads below 5%, and emittances smaller than 2 mm.mrad. A preliminary investigation of the stability of these optimal beams was also carried out. The random scan was enabled by running a large number of fast PIC simulations, which became feasible through the use of several approximations that significantly reduced the computational cost and the access to high-performance Computing (HPC) Irene Joliot Curie facility [4].

The aim of the present study is to use the simulations generated in this fast random scan as training data for machine learning (ML) models. These ML models serve as surrogate models (SM) of the PIC simulations. More specifically, the models are designed to predict the electron beam output parameters—namely energy, charge, energy spread, and emittance—from the laser and plasma input parameters of the PIC simulations, which include a_0, gas pressure, nitrogen concentration, and focal position in vacuum. In this work, we demonstrate the feasibility of using ML methods to construct efficient SM of the PALLAS laser-plasma injector (LPI). The main motivation is the computational efficiency of such models: once properly trained, SM can achieve high accuracy while providing extremely fast inference compared to PIC simulations.

These models can be used to rapidly explore the parameter space and to identify both optimal working points and the stability of different regions of the LPI. They can also be integrated into a comprehensive start-to-end simulation framework for advanced laser-plasma-based electron accelerators. In addition, SM can guide the selection of specific LPI regions where more refined PIC simulations may be performed to study a particular phenomenon, or where targeted experiments can be carried out. Moreover, these models could be integrated with experimental data in order to enable efficient feedback loops for LPI beam stabilisation and control during real-time opera-

tion.

This section begins with a brief overview of the machine learning (ML) techniques used in the following sections. In section 3, we describe the setup and input parameters used for the simulations. Section 4 presents the data sets and discusses the construction of a simple model for predicting injection and filtering data. In section 5, we examine different machine learning approaches for surrogate modelling. Section 6 highlights the performance of these models, discusses several optimisation strategies, and compares surrogate modelling (SM) with conventional methods. Section 7 presents a model of spectrum reconstruction. Section 8 briefly presents work on reverse model. Finally, the concluding remarks summarise the study's impact on LWFA numerical optimisation.

3.2 Machine learning methods

The nonlinear nature of LWFA makes numerical modelling techniques, such as particle-in-cell (PIC) simulations [5], essential for designing reliable laser–plasma accelerators. However, these simulations can become computationally intractable if one relies solely on limited experimental data and scaling laws.

Machine learning (ML) techniques [6] are increasingly employed in LWFA studies and experiments due to their ability to perform nonlinear regression at high speed. Furthermore, frugal models[1] have demonstrated strong predictive performance in numerous nonlinear physics problems [7].

The goal of any ML algorithm is to predict an output y from an input x. To achieve high performance, ML models must be trained using a sufficiently large and representative dataset. Once trained, a model can be used to make inferences on unseen data. A well-trained model generalises effectively to new, previously unseen cases.

We begin by presenting one of the most widely used algorithms: the multilayer perceptron (MLP).

3.2.1 MLP

The MLP is one of the most simple architecture. It consists of several neurons organised in layers. Each layer contains a certain amount of neurons and each neuron inside a given layer receives as input the output of every neurons from the preceding layer. The algorithm then cascades over each layer untill the output is reached. To explain this in a bit more detail, I will go to each step of this process. For an illustration see Fig. 3.1:

[1]Frugal machine learning focuses on developing highly accurate models while minimising data and resource consumption.

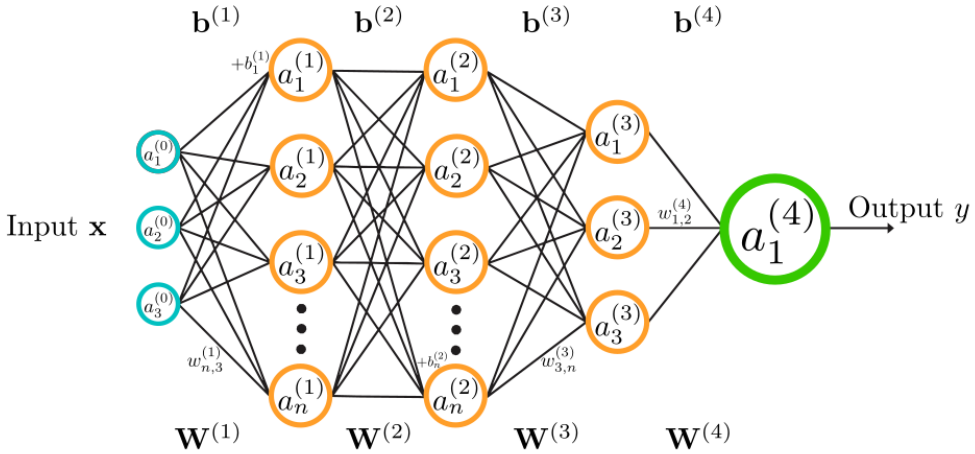

Figure 3.1: Illustration of a MLP with 5 layers, one input layer,3 hidden layers and one output layer. The lines connecting each neuron represent the forward propagation formula (3.1). Figure adapted from apted QuantuMechaniX8, Wikimedia Common

Forward propagation

The first step of training is called forward propagation, the role of this step is to treat the data starting from the input x to arrive at the output y. More formally the forward propagation can be expressed as follow:

$$y_{ik} = \sigma\left(\sum_{j=1}^{N} w_{ijk}x_j + b_{ik}\right) \tag{3.1}$$

In the above sum, y_{ik} represents the i-th neuron on the k-th layer,x_j represents the output of the j-th neuron in the preceding layer, w_{ijk} are the N weights of the neuron y_{ik} associated with every preceding neuron's output, b_{ik} is the bias associated with the neuron y_{ik}, and finally σ represents a nonlinear function (or activation function) that is applied to the sum. To put it in plain English, each neuron is a weighted sum of all the preceding neuron's outputs over which we apply a nonlinear function. One will note the similarity with the expression of a linear regression; in fact, the only difference comes from the nonlinear function σ. The addition of this nonlinear function is of great consequence. Indeed, the Universal Approximation Theorem [8] states that a neural network can approximate any continuous function with an arbitrary degree of accuracy provided that it has a sufficient number of neurons, at least one hidden layer, and a nonlinear function . The most common activation function used in neural networks is ReLU.

$$ReLU(x) = \begin{cases} x & if \quad x > 0 \\ 0 & if \quad x < 0 \end{cases}$$

ReLU is oftern used because it is computationally cheap, facilitate the gradient descent (cf later) and have demonstrated great performances in a wide variety of problems.

Loss function

As we will see in the next section the trainig process of MLP is based on a gradient descent method which aims at minimising a loss function. There is two main classes of loss function supervised losses and unsupervised loss. The supervised loss is often used when the training set consists of N inputs x, that are associated to N outputs y. In this case the loss function is some form of distance metric which measures how close the prediction is to the real data. The most ubiquitous supervised loss is the mean squared error (MSE):

$$MSE = \frac{1}{N}\sum_{j=1}^{N}(y_{j,pred} - y_{j,true})^2 \tag{3.2}$$

Here $y_{j,pred}$ represents the model prediction and $y_{j,true}$ is the value of the target data. An unsupervised loss function, on the other hand, is a function of the input parameters x themselves. This method of learning is used either when the outputs are not available or when the goal is to discover patterns or structure in the data. In either case, the loss function must be chosen appropriately depending on the problem we are trying to solve, since this is what will guide the refinement of the model with the hope of generalisation.

Back propagation

Back propagation is a gradient descent algorithm, the goal of the algorithm is to minimise the loss function at the next iteration of forward propagation. To do so, the algorithm computes the derivative of the loss function with respect to every weights and biases of the model using the chain rule:

$$\frac{\partial L}{\partial w_{ijk}} = \sum_{\text{all paths from } w_{ijk} \text{ to } L} \frac{\partial L}{\partial y^{(M)}} \cdot \frac{\partial y^{(M)}}{\partial y^{(M-1)}} \cdot \frac{\partial y^{(M-1)}}{\partial y^{(M-2)}} \cdots \frac{\partial y^{(k+1)}}{\partial y_{ik}} \cdot \frac{\partial y_{ik}}{\partial w_{ikj}} \tag{3.3}$$

Where L is the loss, y_{ik} is the neuron connected to the weight w_{ijk}. $y^{(M)}$ denotes a neuron in the last layer in a network containing M layers.

During training, the algorithm processes the dataset in iterations called epochs, where one epoch corresponds to a complete pass over all training examples. Within each epoch, the dataset may be divided into batches containing a fraction of the training examples and the gradients are calculated over each batch. The weights and biases are then updated after each batch using the following rule:

$$w_{ijk}(t+1) = w_{ijk}(t) - \lambda \frac{\partial L_{\text{batch}}}{\partial w_{ijk}(t)}$$

where λ is the learning rate controlling the step size, and L_{batch} is the loss computed over the current batch. This is a basic gradient descent update; other optimisation methods such as momentum or Adam [9, 10]may use different update rules.

3.2.2 Gaussian process regression

Gaussian Processes (GP) form a different class of algorithms that, instead of focusing on tuning weights for accurate inference, rely on a fully probabilistic approach. Rather than directly modeling the target function y, a GP defines a probability density over an infinite set of possible functions. This is achieved by relying on the Multivariate Normal distribution (MVND):

$$p(f \mid \mu, \Sigma) = \frac{1}{(2\pi)^{n/2}|\Sigma|^{1/2}} \exp\left[-\frac{1}{2}(f-\mu)^T \Sigma^{-1}(f-\mu)\right], \tag{3.4}$$

where f is the function we aim to model, μ is the mean vector, and $\Sigma = \text{cov}(f)$ is the covariance matrix.

Given a function $f \in \mathbb{R}$ of input parameters $x_i \in \mathbb{R}^m$ with m the dimension of the input and n the number of points at which we evaluate the function, we can represent f as a vector:

$$\vec{f} = [f(x_1), f(x_2), \ldots, f(x_n)]^T.$$

This probability density expresses the probability of observing a particular realisation of f given a mean vector μ and covariance matrix Σ.

The main idea behind Gaussian Processes is to model the function based on a prior belief about its behavior. This belief is encoded in the covariance matrix, which specifies how pairs of inputs are correlated.

A widely used covariance kernel is the Radial Basis Function (RBF) or Gaussian kernel:

$$k(x, x') = \sigma_f^2 \exp\left(-\frac{\|x - x'\|^2}{2l^2}\right). \tag{3.5}$$

Here we have $cov(f(x), f(x')) = k(x, x')$. This kernel encodes the intuition that if two inputs x_1 and x_2 are close to each other in input space, then their corresponding outputs will be highly correlated, whereas inputs far apart will have there correlation go towards zero as illustrated in Fig. 3.2. The parameter l is called the length scale and controls the rate at which correlation decreases. σ_f^2 is the variance. It controls the range of values a function drawn from the Gaussian process can take.

When it comes to training, we need to consider pairs of labelled data (x_i, y_i) with y_i the known output of the function we are modelling at the input points x_i. Using

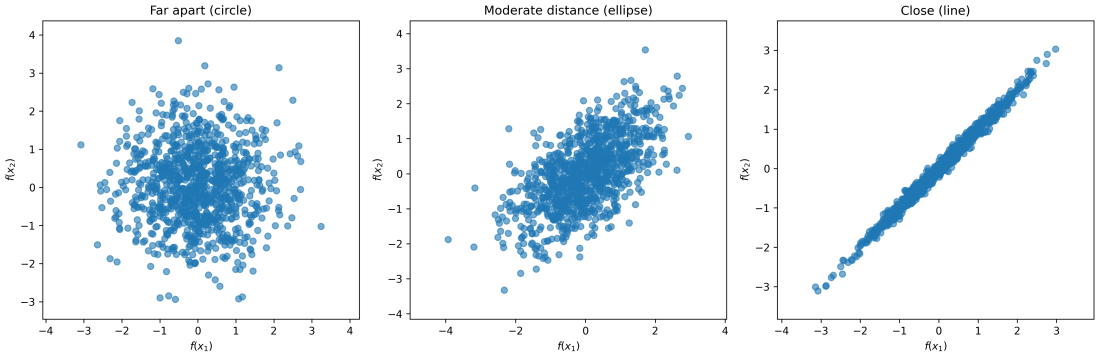

Figure 3.2: Example of three scatter plots of $f(x_1)$ vs $f(x_2)$, with each point representing a different function f drawn from the same Gaussian process with an RBF kernel of lengthscale $l = 1$ and $\sigma_f^2 = 1$. In the left panel, the distance between x_1 and x_2 is 10; in the center, 1; and in the right, 0.1.

this knowledge and the prior distribution (the MVND whith a chosen kernel) we can generate a posterior distribution that encodes the updated knowledge via Bayes' rule. In practical form, letting $X = \{x_1, \ldots, x_n\}$ be the training inputs, $y = [y_1, \ldots, y_n]^\top$ the training targets, and $X_* = \{x_*^{(1)}, \ldots, x_*^{(m)}\}$ the test inputs,i.e the point at which we want to evaluate the modelled function f. Assuming Gaussian observation noise $y = f(X) + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma_n^2 I)$ and a zero mean prior ($\mu \equiv 0$) , the posterior predictive is

$$f(X_*) \mid X, y \;\sim\; \mathcal{N}\big(\mu_*,\, \Sigma_*\big),$$

That is a function drawn from a MVND with updated posterior mean μ_* and covariance Σ_*:

with

$$K_{i,j} = k(x_i, x_j)$$

$$\mu_* = K(X_*, X)\big[K(X, X) + \sigma_n^2 I\big]^{-1} y$$

$$\Sigma_* = K(X_*, X_*) - K(X_*, X)\big[K(X, X) + \sigma_n^2 I\big]^{-1} K(X, X_*).$$

These posterior mean μ_* and covariance Σ_* are what we use for prediction: the posterior mean is commonly taken as the point prediction at X_*, and the diagonal of Σ_* gives predictive variances at that point. In comparison to other machine learning methods, Gaussian process has the advantage to give for prediction both a mean and a variance, which quantifies how uncertain we are about the prediction at a specific point. This quantification of the uncertainty allows such model to be used in optimisation loops such as Bayesian optimisation [11]. Another important thing is that the type of function drawn from a Gaussian process regression are entirely determined by the choice of the kernel and its parameters as illustrated in Fig . 3.3. Lastly the training process only

consists in computing $\left[K(X,X)+\sigma_n^2 I\right]^{-1} y$, which computational cost grows with the cube of the number of training points.

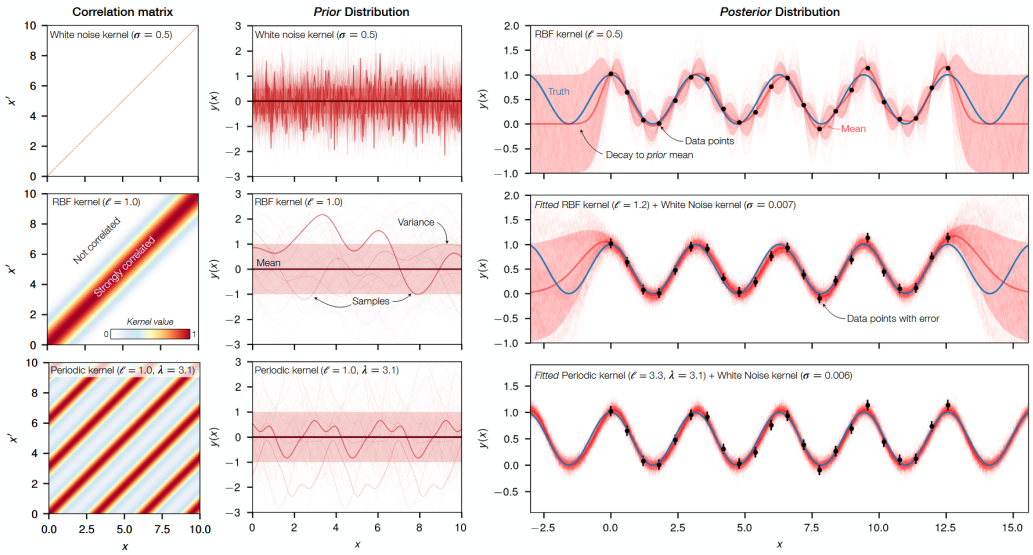

Figure 3.3: The left panels show three different kernels: white noise (top), RBF (middle), and periodic (bottom). The center panels display prior distributions associated with each kernel. For the white noise kernel, functions f drawn from the GP resemble random noise; with the RBF kernel, they appear as smooth continuous functions; and with the periodic kernel, they exhibit periodic behavior. The right panels present the posterior distributions, where the red curve denotes the posterior mean, the black points indicate observed data, and the blue curve represents the target function. The shaded reddish region illustrates the variance, which is largest for the white noise kernel. Since the target function is a cosine, the periodic kernel achieves the best regression performance. Figure adapted from [6]

3.2.3 Decision trees

Decision trees are a machine learning technique that can be used for both classification and regression tasks. The name of this method comes from the tree-like architecture used to make decisions. Indeed, the input data starts at the root node and is recursively split into branches until reaching a leaf node, where the final prediction is made (see Fig. 3.4).

Each split occurs at an intermediate node and consists of a conditional statement (e.g., if $x \leq 20$, take the left branch; otherwise, take the right branch). The main chal-

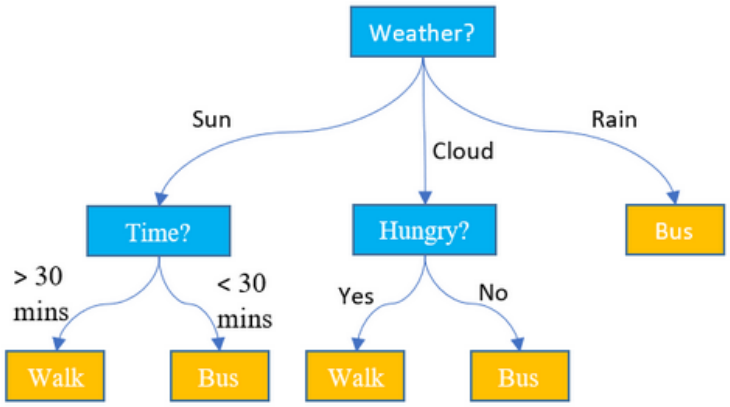

Figure 3.4: Illustration of a classification Decision tree, adapted from, https://www.displayr.com/what-is-a-decision-tree/

lenge in decision trees lies in determining the most appropriate split. This process is performed during training and depends on the specific type of tree algorithm. For example, in Random Forests, the split is chosen by selecting the value that maximally reduces a loss function metric, which depends on the task type (regression or classification).

The splitting process continues until either a pure leaf node is reached,i.e., a node where no further split can improve the prediction (for classification, this corresponds to all training samples in the node being correctly classified), or a pre-set limit specified by the user is reached, such as a maximum depth. The maximum depth is a hyperparameter that governs the maximum number of splits allowed in the tree model.

The main advantages of this method is that they are really fast to train and make inference even though they tend to overfit the data. Overfits is when a model fails to generalise well because they learn a perfect mapping between the training input and training output. These overfits will lead the model to perform badly on unseen data.

3.2.4 Model evaluation

In order to evaluate the performance of a model, several metrics can be used. For instance, in the case of an MLP, one can monitor the loss function over epochs during training. A commonly used performance metric is the coefficient of determination, also known as the R^2 score:

$$R^2 = 1 - \frac{\sum_{j=1}^{N}(y_j - f_j)^2}{\sum_{j=1}^{N}(y_j - \bar{y})^2}. \tag{3.6}$$

Here, y_j denotes the true target value at the j-th sample, f_j is the corresponding

prediction produced by the model, and $\bar{y}$ is the mean of all target values in the dataset. The quantity N represents the total number of data points over which the metric is evaluated. Intuitively, this metric measures how much the predictions deviate from the true values, normalised by the variability of the target data. Put differently, if we consider a scatter plot of predicted values f versus true values y, the closer the points are to forming a straight line, the closer R^2 will be to 1 as illustrated in Fig. 3.5 . An R^2 score of 1 indicates perfect prediction, while a score close to 0 means the model performs no better than simply predicting the mean of the target values. Values of R^2 smaller than 0 are also possible, in which case the model is performing worse than the mean predictor, and in principle R^2 can be arbitrarily negative.

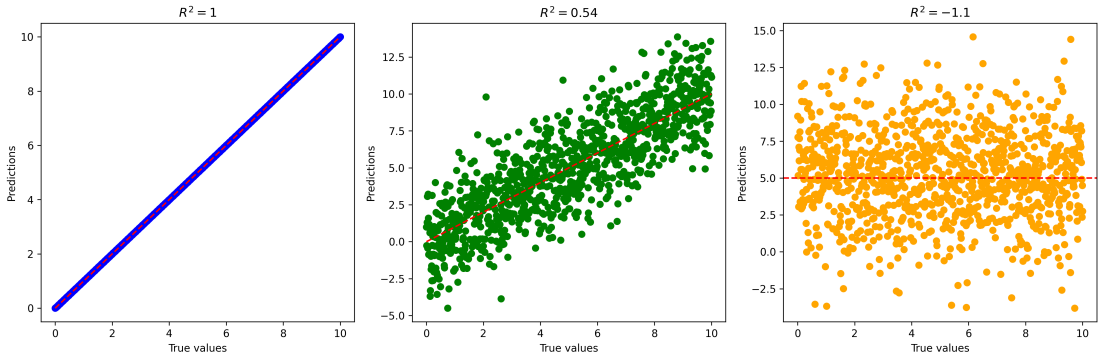

Figure 3.5: Example of 3 scatter plots with varying R^2 score, with $R^2 = 1$,$R^2 = 0.54$,$R^2 = -1.1$ from left to right

3.3 Numerical experiment

The data set used for the SM training comes from PIC simulations aiming to deliver electron beams ranging from $150 - 250$ MeV, $30 - 50$ pC of charge, an energy spread lower than 5% and an emittance of less than 2 mm.mrad as presented in [1]. The LASERIX platform at IJCLab provides the laser driver with a power in the range of 40 to 80 TW. The LPI relies on an ionisation injection scheme [12] with a plasma target divided into two regions [13]. The first region comprises a gas mixture of He doped with N_2 whose length is 0.6 mm. The inner shell electrons of N^{5+} and N^{6+} can be injected in the plasma wakefield. The second region is composed of pure He, 1.2 mm long and dedicated to the acceleration of the injected electrons. The main objectives of the numerical experiments

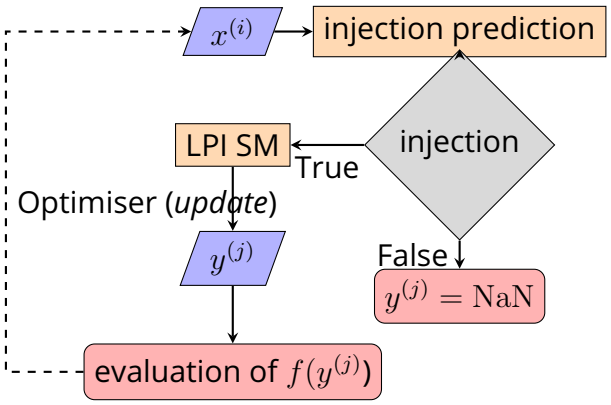

Figure 3.6: Surrogate models and injection prediction-based LPI design optimisation studies flowchart. The dashed line represents the potential optimisation loop update of the input $x^{(i)}$.

are: (i) construct a classification-based model that predicts electron beam injection as a function of input parameters; (ii) construct ML-based SM from simulation data that are able to predict the LPI electron beam parameters; (iii) use both classification model and SM to optimise LPI configurations and investigate the stability of optimal beam parameters.

The scheme in Figure Fig. 3.6 illustrates the principle of the numerical experiments. An LPI configuration input $x^{(i)}$ can be filtered by an injection prediction model, providing inputs for the LPI SM to predict the electron beam parameters $y^{(j)}$. An objective function f built from the outputs $y^{(j)}$ can feed an optimisation routine.

3.4 Data sets generation

As mentioned at the beginning of this section, the main bulk of the simulations used as training data for the surrogate models originated from the fast random scan performed

by Pierre Drobniak et al. In this section, we describe how these simulations were configured and the rationale behind the chosen settings.

Several approximations were applied to reduce the computational cost of the fast PIC simulations. The laser was modeled as a linearly polarised, fifth-order flattened Gaussian beam (cite) with a beam waist of $w_0 = 19\,\mu$m and a full width at half maximum (FWHM) duration of $\tau_0 = 35$ fs.

These values for the beam waist and pulse duration correspond to the optimal configuration of the LASERIX driver. The laser was described using the envelope approximation, which further decreased the computational cost. In addition, in this work only the first mode of an azimuthal mode decomposition was used to model the laser,i.e cylindrical symmetry was assumed thereby further reducing the numerical effort. In the following chapter on realistic simulations, I will employ a greater number of modes and analyse the impact on electron beam quality.

Because the ionisation thresholds of helium electrons and the first five electrons of nitrogen are $9 \times 10^{15}\,\mathrm{W/cm^2}$ and $1.5 \times 10^{16}\,\mathrm{W/cm^2}$ respectively, two to three orders of magnitude lower than the laser peak intensity, these electrons were assumed to be fully ionised at the laser entrance. Consequently, they were treated as a cold plasma fluid in the simulations, eliminating the need to explicitly calculate their ionisation and thus saving computational resources.

The number of macroparticles per cell (MPC) was also reduced from 8 to 1, lowering the computation time from 450 core-hours to 130 core-hours at the GENCI High-performance Computing (HPC) Irene Joliot Curie facility [4]. This reduction had only a modest impact on the simulation results, as reported in [1].

A previous hydrodynamic study conducted with the OpenFOAM software (citer these) demonstrated that the gas cell used in this work exhibited two notable characteristics: first, the transverse pressure profile was nearly uniform at a given longitudinal position; second, the longitudinal profile maintained the same shape over a wide range of pressures. These properties justified approximating the density profile using a polygonal shape. Fig. 3.7 shows a typical longitudinal density profile employed to initialise the simulations. Another important observation was that both simulations and experiments indicated that the nitrogen dopant remained confined when the pressure in the first chamber was approximately equal to that in the second chamber. This observation motivated the choice of using a single pressure as a variable parameter in the simulations, since $p_1 \simeq p_2$ when the dopant was effectively confined.

The extensive approximations used to generate these simulations, enabled a substantial reduction in computational time to 130 core-hours per simulation, corresponding to roughly 30 minutes using our HPC resources settings. This efficiency allowed an initial exploration of the parameter space and the testing of a wide variety of input conditions using a random scan approach.

The simulations involved four input variable parameters, namely $x^{(i)} = (a_0, x_{off}, p_1, c_{N_2})$,

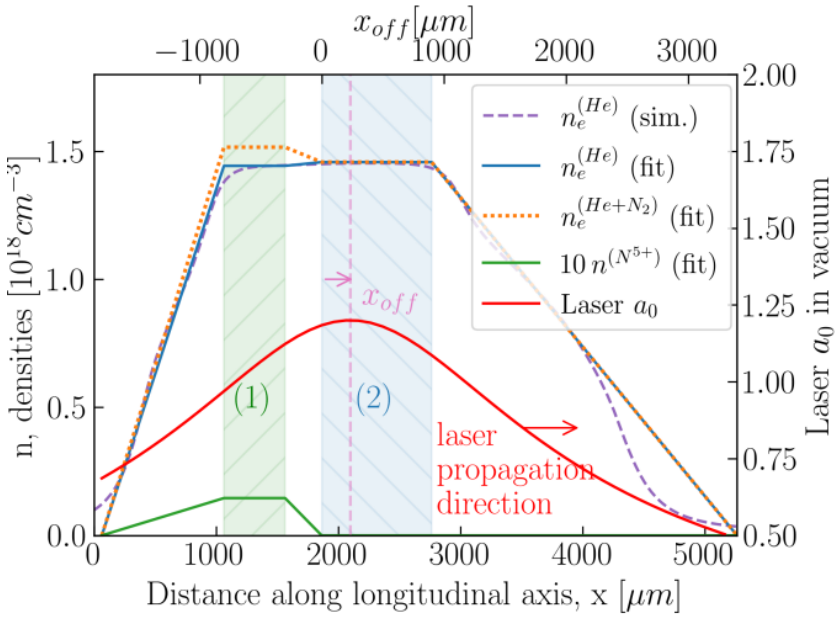

Figure 3.7: Generic longitudinal density profile (blue line) used in the random scan as a fit on the original OpenFOAM profile (purple dashed line). Here the example is a He profile with a plateau at 30 mbar and dopant concentration cN 2 = 1 st stage (green hatched area) and 2 nd stage (blue hatched area) respectively. The dopant density is in green line and the total electron density assuming full ionisation of all species is in yellow dashed line. The envelope of a laser in vacuum with a maximum a0,vac,max = 1.2 reached for x = xof f (focal plane in vacuum) is added. Figure adapted from P.Drobniak phd thesis

where a_0 was the laser pulse normalised vector potential in vacuum, x_{off} the laser focal position in vacuum, p_1 the gas pressure in the first region, and c_{N_2} the concentration of nitrogen in the same region. As depicted in Fig. 3.7, $x_{off} = 0$ corresponded to the entrance of the second chamber [14]. These parameters were selected because they provided a sufficient basis for adjusting and controlling the electron beam parameters in the design of the LPI project. The input parameters $x^{(i)}$ varied within the ranges indicated in Tab. 3.1, which corresponded to the experimentally achievable conditions of the LPI at that time.

a_0	x_{off} [μm]	p_1 [mbar]	c_{N_2} [%]
$[1.1, 1.85]$	$[-400, 1800]$	$[14, 119]$	$[0.2, 12]$

Table 3.1: Interval for the 4 input parameters used for the random scan simulations

The simulations performed by Pierre Drobniak et al. consisted of five massive random scans, each with 2401 configurations. Each random scan explored a part of the input parameter space using either a continuous uniform distribution or a skew-normal distribution. These random scans resulted in some intervals of the input space being over-represented in the data set. The dataset generated by P. Drobniak et al. is refered to as set1.

The density of points within the random scan data set set1 was high enough to explore the input parameter space finely. The resolution is largely above the one that can be obtained experimentally. For example, with x_{off}, we reach a numerical resolution as low as $1\,\mu$m where it is barely $50\,\mu$m in standard experimental conditions.

However in order to improve the performances of the SM as we will see later a second simulation data set set2 was produced using an injection prediction model (cf later) for filtering input parameters $x^{(i)}$ resulting in 3536 simulations. For the set2 data, the input parameters $x^{(i)}$ were randomly drawn from the intervals presented in Tab. 3.1 using a continuous uniform distribution. The histograms in Fig. 3.8 present the configuration distribution of the inputs $x^{(i)}$.

We chose 4 output parameters to characterise the electron beam, namely: $y^{(j)} = (E_{med}, \delta E_{mad}, Q, \epsilon_y)$. E_{med} is the median energy, $\delta E_{mad} = \sigma_{mad}/E_{med}$ with σ_{mad} the median absolute deviation, Q the charge and ϵ_y the transverse normalised emittance. The output parameters $y^{(j)}$ used as validation data in the model were evaluated only at the last time step of the simulation and thus represent the features of the electron beam right after the plasma outramp. They are retrieved using a postprocessing script based on APtools python code [15]

For the set1, out of the 12004 simulations, 9846 resulted in injected beams. The output $y^{(j)}$ ranges are presented in Tab.3.2 for both data set set1 and set2. The two data sets were used both as training data or test data.

data set	E_{med} [MeV]	δE_{mad} [%]	Q [pC]	ϵ_y [mm.mrad]
set1	$[41, 355]$	$[0.01, 58]$	$[3, 791]$	$[0.6, 78]$
set2	$[44, 368]$	$[0.01, 56]$	$[3, 837]$	$[0.6, 76]$

Table 3.2: Interval for the 4 output electron beam parameters of the simulations data set

3.4.1 Injection prediction model

We constructed an injection model trained on the initial data set SET1. This model predicts whether injection will happen for the input $x^{(i)}$. We define the injection as an accelerated beam ($\gamma_z > 10$) with charge $Q > 3$ pC.These thresholds were selected to exclude low-energy particles and numerically insignificant bunches, while retaining beams that

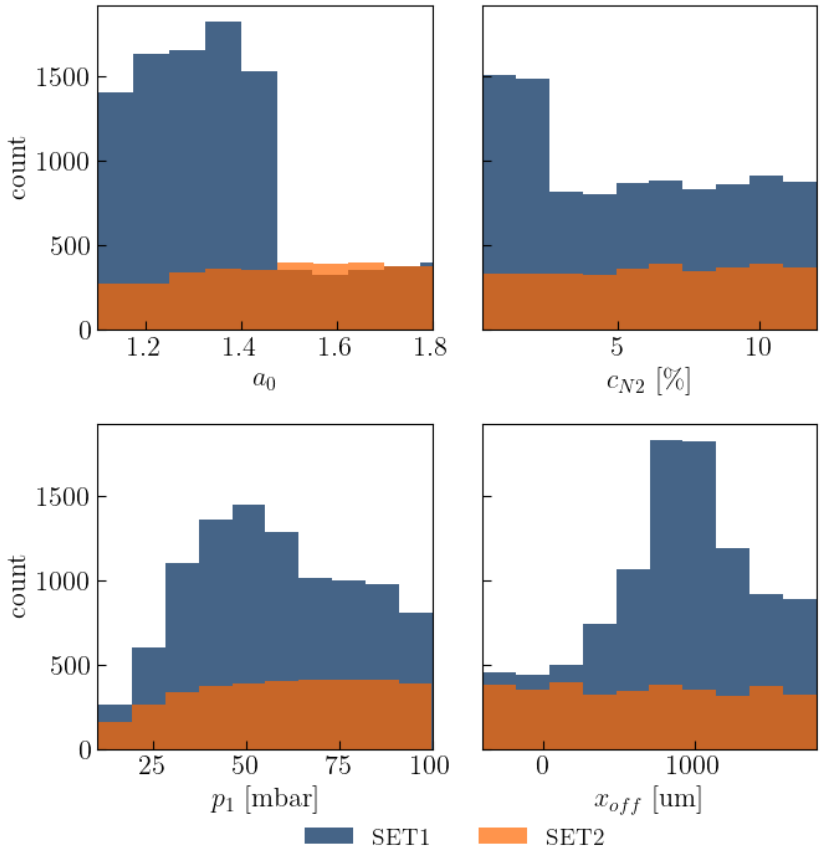

Figure 3.8: Distribution of the input parameters for the 9846 training data simulations (blue) and the 3536 test data simulations (orange).

would be relevant for LPI operation. The injection model uses a simple random forest algorithm to make its prediction. The accuracy of the model is 98%. Accuracy denotes the number of correct predictions over the total number of predictions. This injection prediction model can save computational time before launching new simulations by avoiding to compute simulations that would result in beams below the injection threshold. It can also be used as a constraint in the Bayesian optimisation search to filter the input parameters.

Fig. 3.9 shows, as a function of the input parameter (a_0,p_1): in blue, the points corresponding to no injection, in orange the points corresponding to injection. The black curve corresponds to the critical pressure p_c for self-focusing as defined in [1, 16].

$P_c = 17 \left(\frac{\omega_0}{\omega_p}\right)^2 [GW]$

and using the following definitions:

$w_L = w_0 \sqrt{1 + (\frac{z}{z_r})^2}$

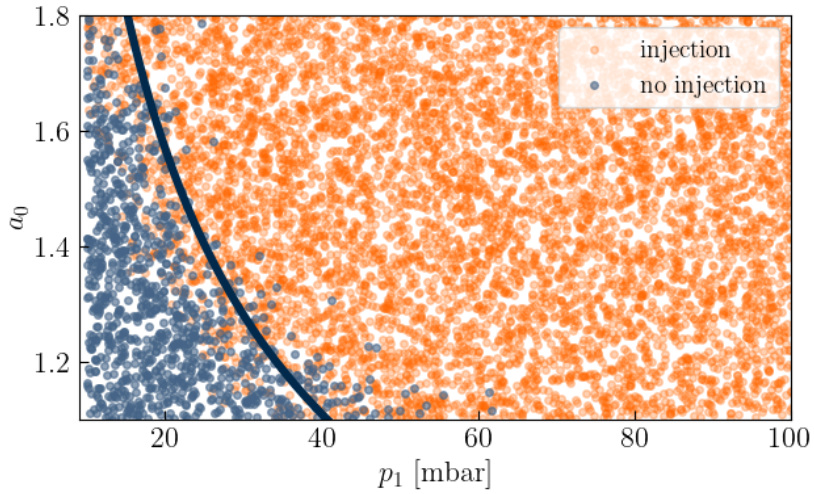

Figure 3.9: Injection model tested on randomly generated points. The dark-blue curve is $3 \cdot p_c$ with p_c the critical pressure necessary for self-focusing which is calculated using the mean value of c_{N2} and x_{off} over the all data set.

$\omega_p^2 = \frac{n_e e^2}{m_e \epsilon_0}$;
$p_{tot} = p_{He} + p_{N2} = n k_B T$
$p_{He} = \frac{n_{eHe}}{2} k_B T = (1 - C_{N2}) p_{tot}$
$p_{N2} = \frac{n_{eN2}}{10} k_B T = C_{N2} p_{tot}$
$n_e = n_{eHe} + n_{eN2}$
$p_{tot} = \frac{n_e k_B T}{2(1+4C_{N2})}$
$a_{0eff} = \frac{a_{0max}}{\sqrt{1+(\frac{z}{z_r})^2}}$;$a_{0max} = a_0$ used as simulation input
$a_{0eff} = 0.86 \lambda_0 [\mu m] \sqrt{\frac{2P_c}{\pi w_L^2}[10^{18} W/cm^2]}$
$a_{0eff}^2 = 0.86^2 \lambda_0^2 . 10^{12} \frac{2.17 \omega_0^2 . 10^9 m_e \epsilon_0 . 10^{-18}}{\pi w_L^2 10^4 n_e e^2}$
$a_{0eff}^2 = 0.86^2 . 10^{-1} \frac{136 \pi m_e \epsilon_0 c^2}{w_L^2 n_e e^2}$

$$p_c \simeq 16 \cdot \frac{\epsilon_0 k_B T m_e c^2}{e^2 w_0^2} \cdot \frac{1}{a_0^2} \cdot \frac{1}{(1 + 4 c_{N_2})} \tag{3.7}$$

with w_0 the laser waist, a_0 the normalised vector potential at the position $x_{off} = 0$ in vacuum used as simulation input. It can be seen that the injection condition determined by the model is proportional to the limit set by the self-focusing threshold. Since, the electron injection by ionisation from N^{5+} and N^{6+} limit is greater than the critical pressure for self-focusing, given the limited intensity assumed for the laser driver pulse in the study, $a_0 \in [1.1, 1.85]$. The injection model improves the prediction from simply using the formula with the proper proportionality factor from 95% to 98%

3.5 Model construction

The current section intends to present the training process of the models and the importance of the input parameter distribution in the training data set.

Multilayer perceptron (MLP)

The first method used to construct a surrogate model of the LPI consisted of four different MLP's. These MLP's were implemented using *Tensorflow* and *Keras* python library [17]. Each MLP predicts one output parameter. For E_{med}, Q and ϵ_y. The MLPs have the following architecture: 5 layers in total, 1 input layer with 4 neurons, 1 output layer with 1 neurons, and 3 intermediate layers with 100 neurons. Each layer has a 20% dropout rate. We used the function PRELU [18] as an activation function for the 3 intermediate layers and a sigmoid function for the last layer. Altogether, this model contains 21101 trainable parameters. This model was trained on 200 epochs with a batch size of 50. The loss function used was the mean squared error (MSE). To avoid over-fitting, we also used a *K-fold* cross-validation method [19]. For predicting δE_{mad} the architecture is modified to improve accuracy, since this parameter is highly correlated with energy and charge. The input layer contains 6 neurons instead of 4. These additional input neurons are the prediction of E_{med} and Q from the already trained MLP.

Extreme gradient boosting (XGB)

We implemented this method by using the *xgboost* library[20]. The maximum tree depth was set to 10, and the loss function was also MSE in this case. We used *K-fold* cross-validation.

Gaussian process (GP)

We implemented this method by using the *Scikit-learn* library [21]. The kernel used was *Matérn*, which is a generalisation of the Gaussian radial basis function and allows the capture of physical processes due to double differentiability by the choice of a smoothness parameter $\nu = 2.5$ and a lengthscale $l = 10^{-5}$.

It is important to note that for all models (MLP, XGB, GP), it is mandatory to rescale the outputs and the inputs from 0 to 1 to get the most accurate results. The output parameters, $y^{(i)}$, are scaled so that the calculation of the loss function is well-weighted, corresponding to the same magnitude in all of the outputs.

3.5.1 Training surrogate model construction resources requirement

The training process on a high-performance laptop is relatively quick, and it takes only a few seconds for the XGB model and a few minutes for the GP and MLP models. Additionally, the computation time for LPI configurations is significantly shorter than that of low-fidelity simulations on HPC CPU nodes. The MLP, XGB, and GP models are approximately 10^7 , 10^8, 10^7 times faster than simulations, respectively. The training of the LPI surrogate models, as the number of input and output is limited, was done using a standard I7 Intel cpu.

The following table tab summarises the SM training time and computing time to generate 10^5 LPI configurations.

time [s]	training	computing 10^5 configurations
MLP	850	10
GP	52	14
XGB	13	0.34

Table 3.3: Training and computing time for each type of method used to build the LPI model using I7 Intel cpu.

Importance of the input parameters distribution

All three models MLP, XGB, and GP were tested on the set2 data, consisting of 3536 test points, separate from the 9846 samples of set1 used for training. We observe that the coefficient of determination R^2 between the SM predictions and the outputs of set2 is above 0.85. However, this score significantly decreases in regions where the density of training points is lower than 1. For each test point, the local density is defined as the number of training samples contained within a hypercube of side length 0.1 in the normalized input space, centered on that test point

We propose a reliability criteria for the SM based on the relation between R^2 and MSE :

$$R^2 = 1 - \frac{\sum_{i=1}^{4}\sum_{j=1}^{N}(y_{ij} - f_{ij})^2}{\sum_{i=1}^{4}\sum_{j=1}^{N}(y_{ij} - \bar{y}_i)^2} = 1 - \frac{4\sum_{j=1}^{N} MSE_j}{\sum_{i=1}^{4}\sum_{j=1}^{N}(y_{ij} - \bar{y}_i)^2} \tag{3.8}$$

With y: output value of the test data, f: predicted value by the surrogate, $\bar{y}$: simulation mean value for a batch of size N. The index i represents the 4 output parameters,

and j represents the test points. MSE_j is the mean squared error for a specific test point $MSE_j = \sum_{i=1}^{4}(y_{ij} - f_{ij})^2/4$.

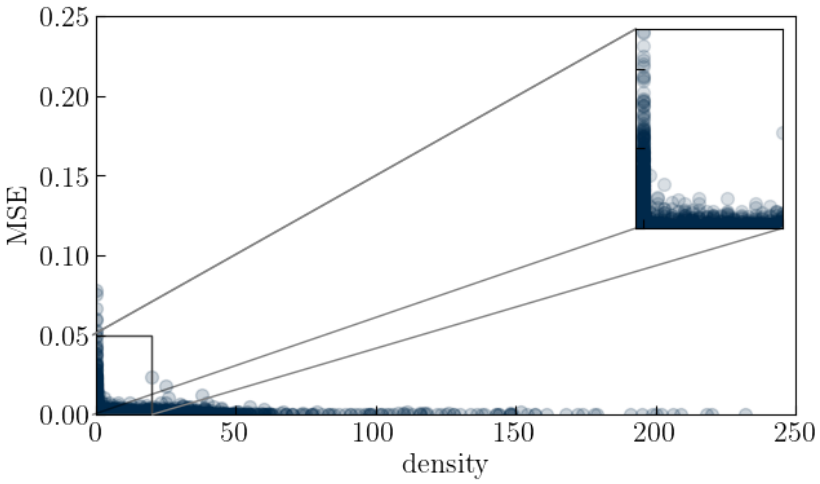

Figure 3.10: Scatter plot showing the relationship between the density of input data configuration points and MSE for every test point for the MLP model. The blue points represent the different configurations.

The relationship between MSE and R^2 can be used to define reliability criteria for the test points in a given region of the input parameter space. We observe that $R^2 \geq 0.9$ for a given batch corresponds to $MSE < 4.10^{-3}$ and that the probability of getting a small MSE increases with the local density of training points as shown in Fig. 3.10. Thus, to ensure that $R^2 \geq 0.9$ in every region of the 4-dimensional input space, a sufficiently high local density is required. For a hypercube of side 0.1, this corresponds to a density ≥ 1. However, when using the simulation data of SET1 [1] as training, this condition was not met in at least 30% of the parameter space. For this reason, the following models were trained on the uniformly distributed data from SET2 and tested on the data of SET1. As shown in Fig. 3.11, 3.12 a better homogeneity largely compensates for reducing the number of training points for our LPI SM.

3.6 Results

3.6.1 Performances of surrogate models

The SM trained on the data set2 showed good correlation with MLP, XGB and GP model having an R^2 score of 0.97, 0.90 and 0.96, respectively, across all output parameters. However, as shown in Table 3.4, the median energy and charge are consistently better

predicted by the models compared to δE_{mad} and ϵ_y. To compare the results of the SM with a more standard method, we added nearest-neighbour interpolation.

	E_{med}	δE_{mad}	Q	ϵ_y
MLP	0.99	0.96	0.99	0.95
XGB	0.97	0.88	0.96	0.78
GP	0.99	0.95	0.99	0.90
interpolation	0.90	0.83	0.91	0.72

Table 3.4: R^2 correlation score for the different surrogate models trained on SET2 and evaluated on SET1.

Figure 3.11 illustrates that the MLP and GP model reach $R^2 = 0.93$ with a training size of approximately 500 samples. All SM outperform a simple interpolation model, which only reaches a R^2 score of 0.85 when the training size exceeds 3000 samples, compared to less than 300 samples for MLP and GP SM.

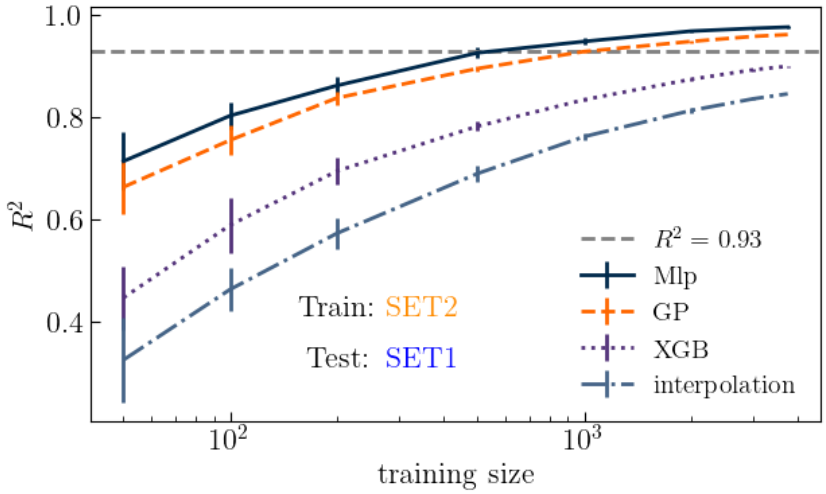

Figure 3.11: coefficient of determination R^2 as a function of the training size in log scale, for all the SM trained on SET2 and tested on SET1. R^2 was taken as the average over 10 training sessions, with the vertical bars representing the standard deviation

To evaluate the performance of the SM across the entire output interval, we computed the mean absolute error (MAE) for each small slice of these intervals.
Fig. 3.13 shows, that the MLP model is the best, followed by the GP. Although the MLP model is the best overall, Fig. 3.14 shows that the GP model is the best in the ranges of interest ($150 - 250$ MeV, with $30 - 50$ pC of charge, an energy spread lower than 5% and

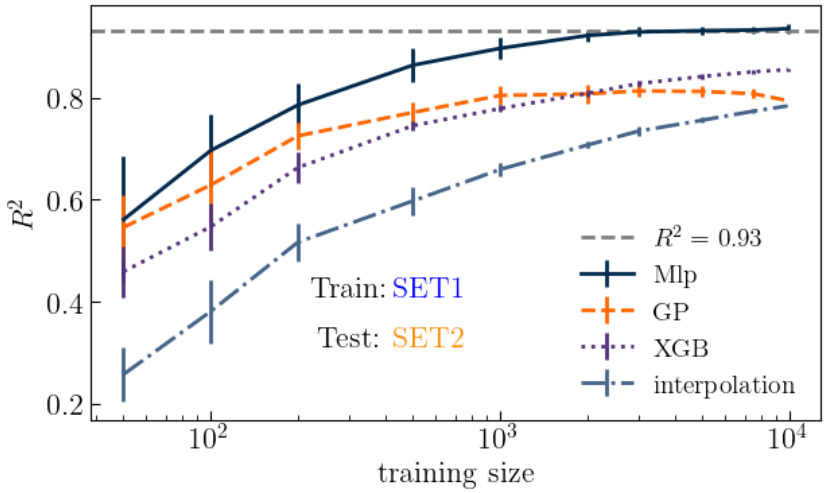

Figure 3.12: Coefficient of determination R^2 as a function of the training size in log scale, for all the SM trained on SET1 and tested on SET2. R^2 was taken as the average over 10 training sessions, with the vertical bars representing the standard deviation.

an emittance of less than 2 mm.mrad). Additionally, the MAE tends to increase for the highest output values since these values are underrepresented in the training data set as shown by the histograms depicting the distribution of the output parameters of set2 in Fig. 3.13.
In Fig. 3.15, the prediction of each SM is represented in a 2D subspace of c_{N_2} and p_1. The projections are made for the input laser parameters fixed to $a_0 = 1.43$ and $x_{off} = -265\,\mu$m. One should notice that a complete set of projections can be generated in a few seconds on a laptop for a complete scan of a_0 and x_{off} or other target parameters for more complex studies.

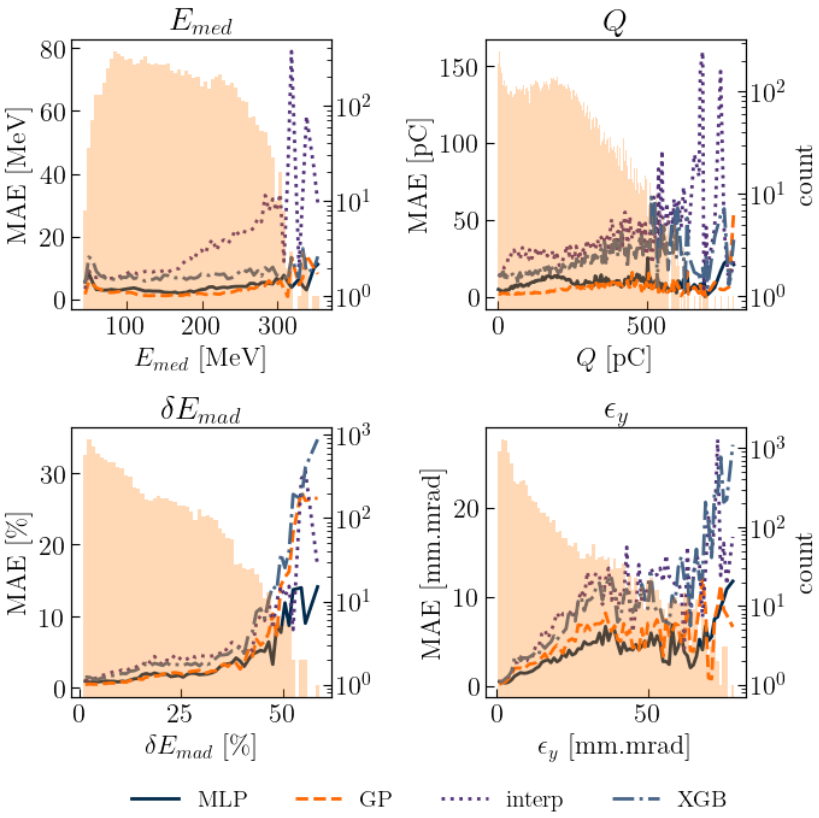

Figure 3.13: MAE for all the SM and all of the output with histograms representing the distribution of the output parameters of training data set set2 in log scale

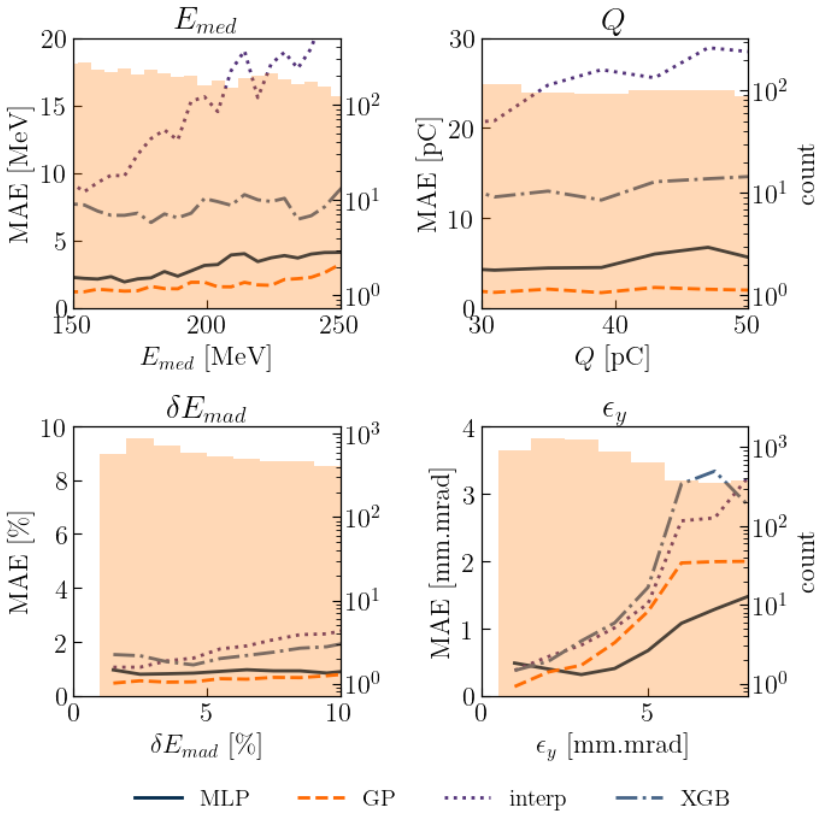

Figure 3.14: MAE for all the SM and all of the output with histograms representing the distribution of the output parameters of training data set set2 in log scale zoomed on the region of interest for outputs.

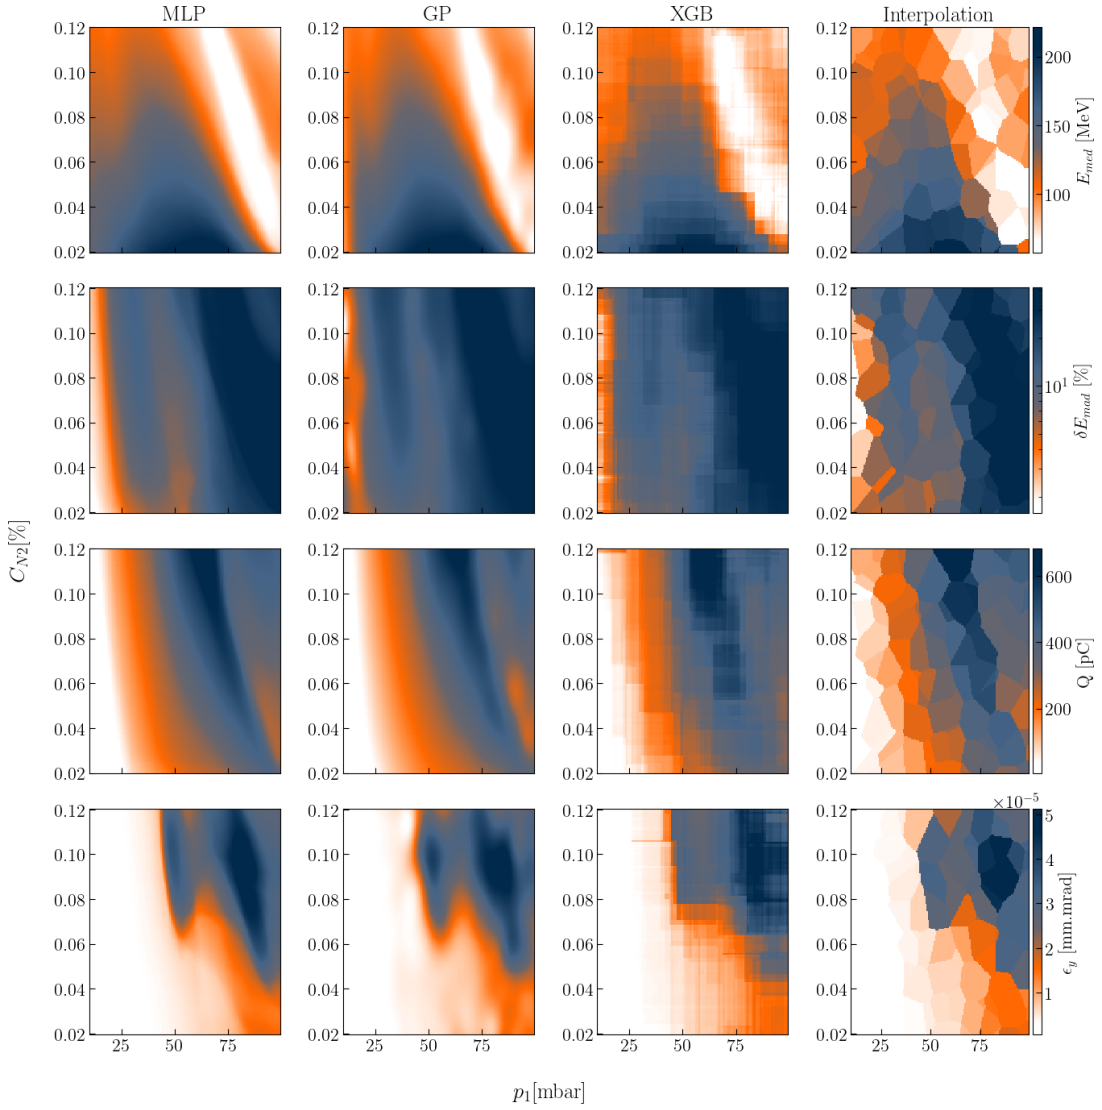

Figure 3.15: Surrogate LPI models prediction for all of the output in a 2D subspace of p_1 and c_{N2} for $a_0 = 1.43$ and $x_{off} = -265\,\mu m$. Other snapshots of 2D subspace can be generated using the python notebook available on the online repository[22]

3.6.2 Optimisation with the surrogate models

The GP model was employed in the subsequent analysis due to its superior performance compared to the XGB and interpolation methods. Although it is less accurate than the MLP across the entire output space, the GP model demonstrates higher efficiency within the specific range of interest $150^{\sim}250$ MeV peak energy, $30^{\sim}50$ pC charge, energy spread below 5%, and emittance under 2 mm·mrad—as illustrated in Fig. 3.14.

Optimum LPI working point stability

Using the GP model, we looked for optimal working points. Several methods can determine these configurations of target and laser for the optimal electron beam parameters. The simplest approach is to generate many data points using a continuous uniform distribution of 4D input parameters. The input range is kept within the boundaries of Tab. 3.1.

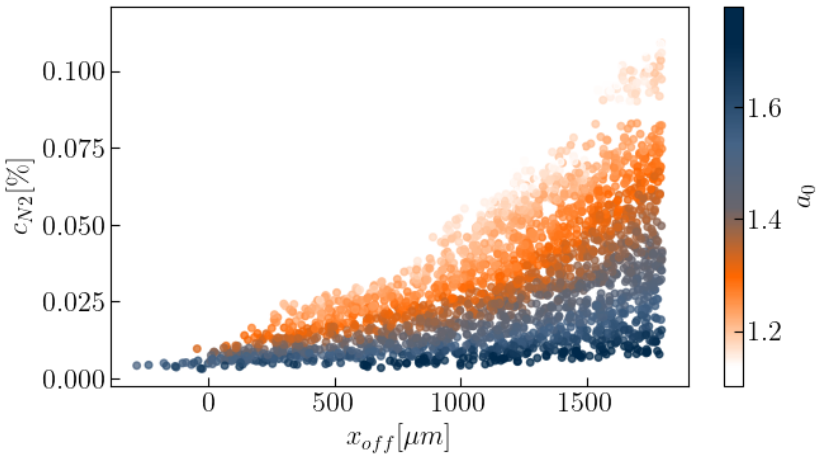

Figure 3.16: 2D subspace of x_{off} and c_{N2} of the input parameters for the configurations selected by filter $\tilde{F}_1$ with a_0 as a color scale.

From this data set, our model can then be used to select beams with the desired characteristics. Selection is performed using the following filter: $\tilde{F}_1 = \{E_{med} \in [205, 215]$ MeV, $\delta E_{mad} < 3.5\,\%$, $Q \in [25, 35]$ pC, $\epsilon_y <$ 2mm.mrad$\}$. We generated 5 million random configurations using a uniform distribution and then used the injection model to keep only the configurations that predict injection. From these configurations, 2347 were selected by filter $\tilde{F}_1$. We can see in Fig. 3.16, for each value of x_{off}, c_{N2} logically decreases with a_0 since the target charge is fixed in the filter. If the laser energy is lower, the charge can be maintained at a certain level by increasing the doping rate, as explained in [1]. Not only the number of injected electrons can be increased, but also increasing the self-focusing helps to reach the threshold value for a_0.

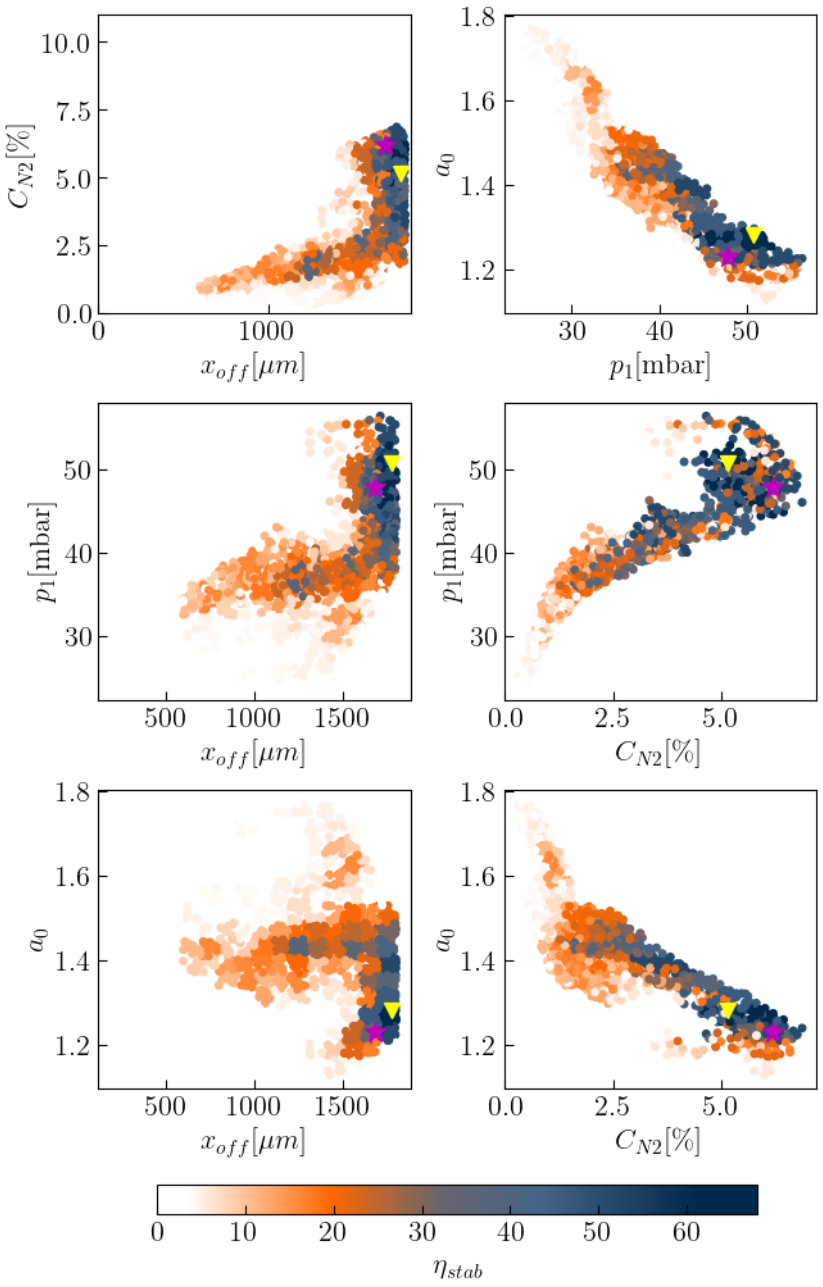

Figure 3.17: Stability map: Projection in the 2D sub-spaces of the configurations selected by filter F_1 with η_{stab} as a color scale. These maps allow us to see the location of the most stable regions. The $\star$ symbol is the configuration 7516 from SET1 and the $\triangledown$ symbol is most stable point in the filter F_1

One interesting aspect to examine with this method is the stability across a target working point. Since the 5 million points were generated using a uniform distribution, each region of the 4D input space contains roughly the same number of points. Thus, we consider that the region with the highest density of remaining points after applying $\tilde{F_1}$ in the input space is the most stable. In Fig. 3.17, we present stability maps as projections in 2D sub-spaces, showing the density of points η_{stab} in the input space of filter $\tilde{F_1}$, the density is represented with the colours scale. These stability maps can guide the search for ideal electron beams. From this analysis, we identified that the most stable region is centred around the following point: $a_0 = 1.31$, $c_{N_2} = 6.1\,\%$, $x_{off} = 1.676\,\text{mm}$, $p_1 = 39\,\text{mbar}$.

Bayesian optimisation

Bayesian optimisation can also be used with SM. We can employ either single or multi-objective Bayesian optimisation (MOBO). For single-objective optimisation, we used the following function to be optimised: $\tilde{f_3} = \sqrt{Q}E_{med}/\delta E_{mad}$ [23]. The optimisation consisted in one hundred steps with 20 random evaluations. Out of the 120 points from the Bayesian optimisation, 39 met or exceeded $95\,\%$ of the maximum of the objective function. These configurations have an average charge of 149 ± 6 pC, an energy of 236 ± 3 MeV and energy spread of $2.8 \pm 0.05\,\%$.

The MOBO aims at optimising simultaneously the elements of the following vector $\tilde{G} = (\delta E_{mad}(x), |E_{med}(x) - E_0|, Q(x))$ where x is the 4D input vector. Our goal here is to maximise charge and minimise energy spread for a given median energy. We tried three MOBO searches with the vector $\tilde{G}$ for three different central median energies 150, 200 and 250 MeV within a ±10 MeV. Each MOBO consisted in 80 steps with 10 random evaluation and 10 evaluation for each step. This MOBO search resulted in Pareto fronts [24], illustrated in Fig. 3.18.

Here, we can see a clear trade-off between charge and energy spread. These solutions correspond to *Pareto*-optima [25] as we cannot improve one of the objectives without deteriorating the other.

These three approaches permit the finding of target working points for the LPI. However, we want to emphasise that the first one with filter $\tilde{F_1}$ is the most mature one because this allows us to find the beams of interest and the most stable LPI configurations. This approach, however, requires generating a large amount of data, which is not a problem with the MLP model contrary to PIC simulations.

3.6.3 Comparison with random scan optimisation

The previous method can be used to seek configurations that outperform the best ones identified in [1]. In this work, two criteria were used for beam selection: $f_3 = E_{med} \cdot Q/\sigma_{mad}$ and filter $F_2 = \{E_{med} > 150\,\text{MeV}, \delta E_{mad} < 5\,\%, Q > 30\,\text{pC}, \epsilon_y < 2\mu m\}$. The

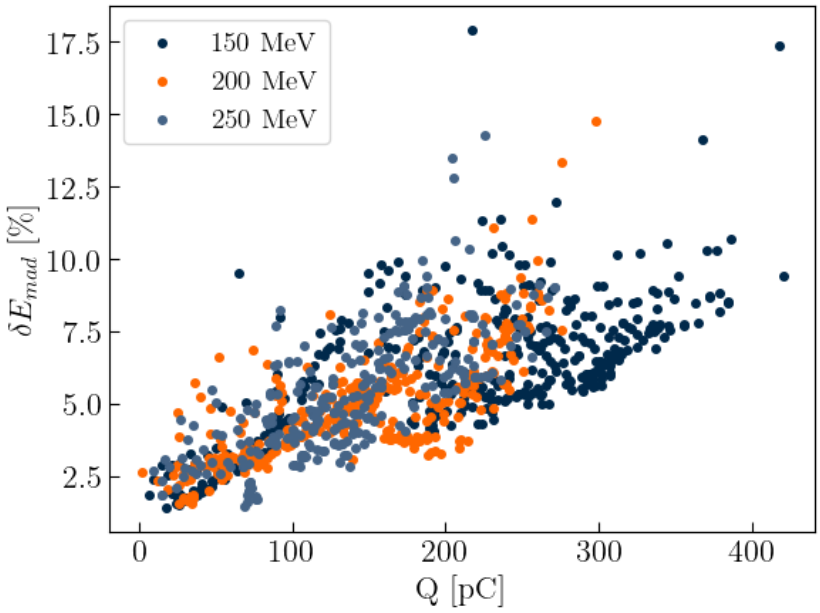

Figure 3.18: Scatter plot showing the result of MOBO search for 3 different energies $150, 200$ and 250 MeV within a ± 10 MeV

best beams were chosen based on the following: The configuration that maximised f_3 (configuration 3702) from SET1, the configuration within filter F that had the smallest δE_{mad} (configuration 7516) from SET1. The characteristics of configurations 3702 and 7516 are presented in Tab .3.5. To compare our method with the results from [1], within the 5 million random configurations, we identified 3779 configurations with higher f_3 values than configuration 3702 and 408 configurations that satisfied the conditions of filter F with a smaller δE_{mad} than configuration 7516. The results are presented in Tab .3.5.

For configuration SM $f_3^{(opt)}$, we have an upstream focus coupled with a high pressure in target chamber 1 and high dopant concentration, leading to strong self-focusing; all of which leads to a high charge beam even with relatively low intensity. However, the high amount of charge leads to a high emittance value. For SM $F_2^{(opt)}$, we have a downstream focus with moderate pressure and intensity, which leads to a low injected charge where very low energy spread is possible. We can thus see that the SM can find working points that outperform a simple random scan. In addition, it is interesting to look at regions of interest, for example, we show in Fig. 3.19. the 2D sub-spaces of all the points that have f_3 values superior to configuration 3702 and identify regions of interest with different colours. At first glance, the different beams are clustered in different areas. A first group of interest is $E_{med} > 215, \delta E_{mad} < 3.53\%$. Most of this set is located at high x_{off} and high p_1 values, with x_{off} above 1500 μm. The combination of high pressure and down-

	$f_3^{(opt)}$ [1]	$F_2^{(opt)}$ [1]	SM $f_3^{(opt)}$	SM $F_2^{(opt)}$
x_offf	558	1680	-372	1798
a_0	1.43	1.23	1.24	1.33
C_{N2}	1.88	6.17	9	7.70
p_1	58.6	47.8	88	40.7
E_{med}	215	212	103	185
δE_{mad}	3.53	1.55	3.09	0.9
Q	198	30	311	45
ϵ_y	5.03	1.74	38	1.5

Table 3.5: Input and output parameters of the best configurations found by the random scan and filter

stream focus leads to a more important a_0 in the accelerating region, leading to a higher wakefield amplitude and high energy beams. This set also displays low c_{N2} and low a_0, which limits the amount of injected charge to reasonable values (58 to 94 pC) despite the high pressure. The reasonable electron bunch charge also allows for maintaining small energy spread by limiting space charge and unwarranted beam loading. The most prominent group is the high charge case $Q > 198\,pC$. This set is distributed all over the hyperspace but we find that it generally follows this guideline the higher a_0, C_{N2}, p_1 the higher the charge will be. Increasing any of these three parameters means increasing the number of inner-shell ionised electron, all other things being equal. The lower x_{off}, the higher the charge will be following the trends of a lower x_{off} balances a higher a_0 in the injection region and thus a higher inner-shell ionisation rate.

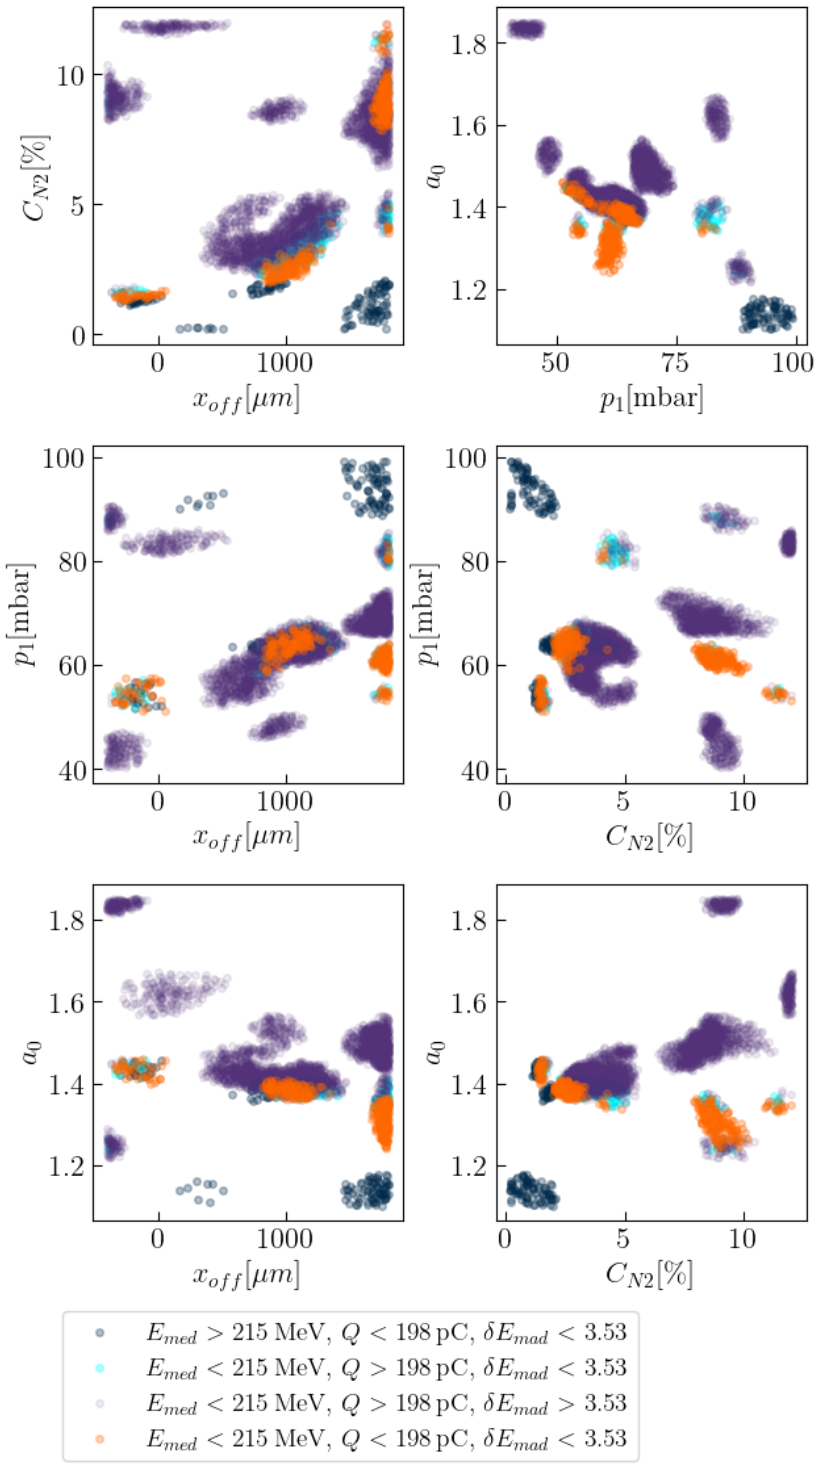

Figure 3.19: 2D subspaces of the configurations selected by the function f_3 higher than 3702 with regions of interest highlighted by different colours.

3.7 Spectrum prediction

In the preceding section, we presented SM trained to predict scalar quantities of the electron beam obtained from PIC simulations. We now extend this approach to predicting the full electron energy spectrum from the same input parameters. Retrieving the 1D spectrum is more informative than recovering only statistical values such as E_{med}, Q, or δE_{mad}.

First, experimental spectra rarely follow simple Gaussian shapes: they often exhibit secondary peaks, long high-energy tails, or abrupt cutoffs. Such features are entirely hidden if one only considers scalar metrics. Second, the shape of the spectrum encodes physical insight for example, a sharp peak can indicate localised ionisation injection, while a plateau may reveal continuous self-injection. Furthermore, statistical values alone can be biased by outliers; for instance, a small fraction of high-energy electrons may skew E_{med}.

Finally, using the spectrum directly allows for application-oriented comparisons: a narrow, quasi-monoenergetic spectrum is far more relevant for many applications than E_{med} or δE_{mad}, which can be misleading when dominated by outliers. Such surrogate predictions could therefore be used to compare directly with experimental spectra, rather than relying only on derived scalar quantities.

In order to construct a surrogate model mapping $[x_{off}, a_0, c_{N2}, p_1] \rightarrow \frac{dQ}{dE}(E)$, we used as input both the variable parameters $x^{(i)} = (x_{off}, a_0, c_{N_2}, p_1,)$ from set1 and set2, as well as the predictions of an already trained surrogate model, namely $SM(x^{(i)}) = (E_{med}, Q, \delta E_{mad})$. In other words, the input space consists of seven parameters: the four variable parameters described above, and three additional parameters obtained from the predictions of a previously trained surrogate model.

The output of this model is the electron spectrum obtained from the post-processing of the SMILEI simulations of set1 and set2. In our post-processing, each spectrum was constructed using 200 bins, corresponding to 200 energy points. To reduce the dimensionality of the problem $\mathbb{R}^7 \rightarrow \mathbb{R}^{200}$,which would otherwise require a very large neural network to correctly learn all 200 targets,we first fitted each spectrum using a Gaussian mixture model (GMM) (cite). A GMM is a statistical model that reconstructs a complex distribution function as a combination of simpler components, in our case Gaussians. The fit obtained with the GMM is extremely accurate, as illustrated in Fig. 3.20. We settled on 7 Gaussians, which already yields a very good reconstruction of the spectra, as shown in the figure.

This approach reduced the dimensionality of the problem from $\mathbb{R}^7 \rightarrow \mathbb{R}^{200}$ to $\mathbb{R}^7 \rightarrow \mathbb{R}^{21}$. Indeed, since the problem involves 7 Gaussians and each Gaussian is fully characterised by three parameters,amplitude, variance, and mean,the surrogate model needs to predict 21 parameters.

In the preceding section, we showed that the MLP provided the best predictions,

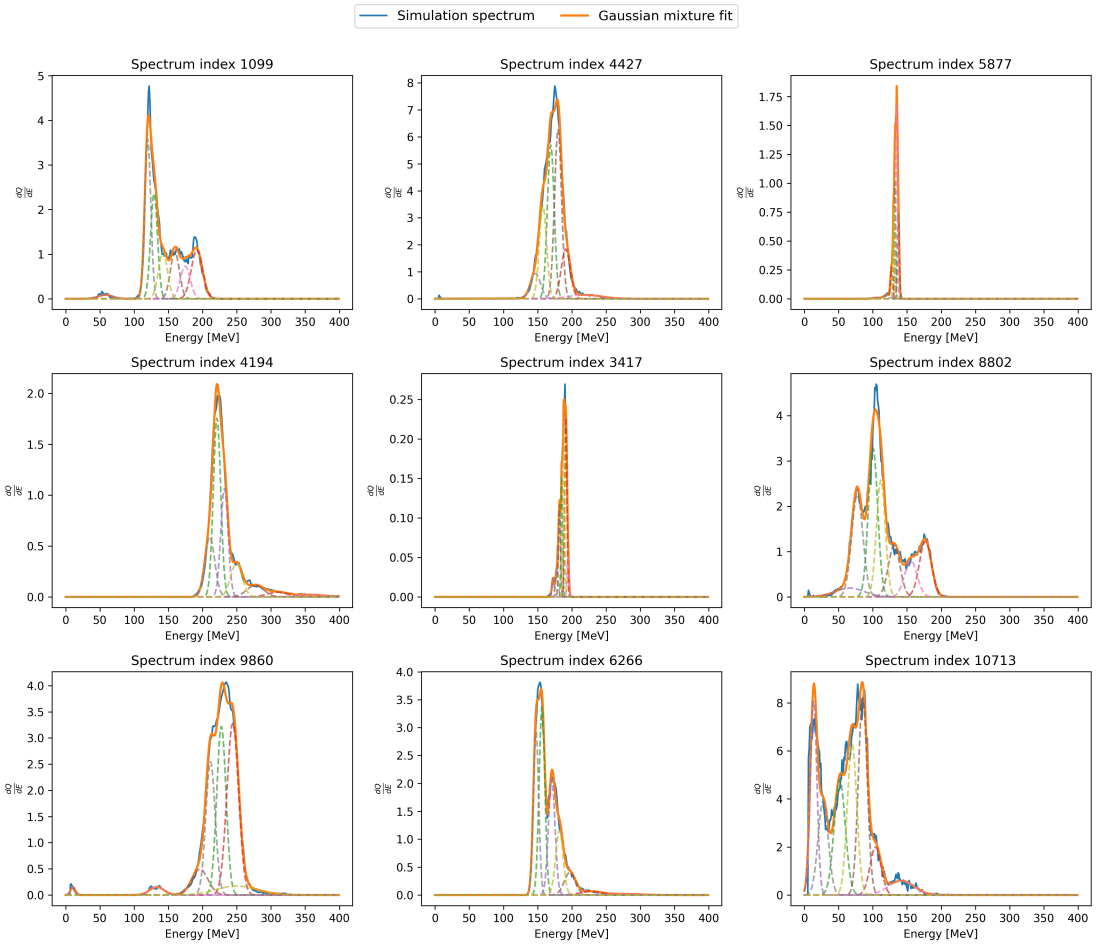

Figure 3.20: An example of 9 randomly chosen spectra fitted using a Gaussian mixture model. The spectra are shown in blue, and the Gaussian mixture fit is shown in orange. The Gaussian components are represented with dashed coloured lines.

and therefore we chose to use the MLP for this task.

To predict the full electron spectrum, we constructed a single MLP model using *Tensorflow* and *Keras* [17]. The input layer contains 7 neurons: the four variable parameters $[x_{off}, a_0, c_{N_2}, p_1]$ and three additional parameters obtained from the predictions of a previously trained surrogate model ($E_{med}, Q, \delta E_{mad}$). The network has three intermediate dense layers with 64 neurons each, using the PRELU activation function [18], and a 20% dropout applied after each intermediate layer. The output layer has 21 neurons, corresponding to the parameters of the 7-Gaussian mixture model (amplitude, mean, variance for each Gaussian), with a linear activation and a non-negativity constraint on the kernel weights.

The model was trained using the Adam optimizer with a learning rate of 10^{-3} and the mean squared error (MSE) as the loss function. To avoid overfitting, a 10-fold cross-validation was applied. A learning rate scheduler and model checkpointing were employed to save the best model for each fold. The training was performed over 200 epochs with a batch size of 50.

This architecture allows the MLP to map the 7 input features to the 21 parameters of the Gaussian mixture, effectively reconstructing the full electron energy spectrum while leveraging previously predicted scalar quantities for improved accuracy.

The MLP-Spec provides excellent reconstruction of the spectra, as illustrated in Fig. 3.21.

The model achieved reasonable results, as illustrated in Figure 3.21, which presents nine randomly chosen spectra and the corresponding MLP predictions.

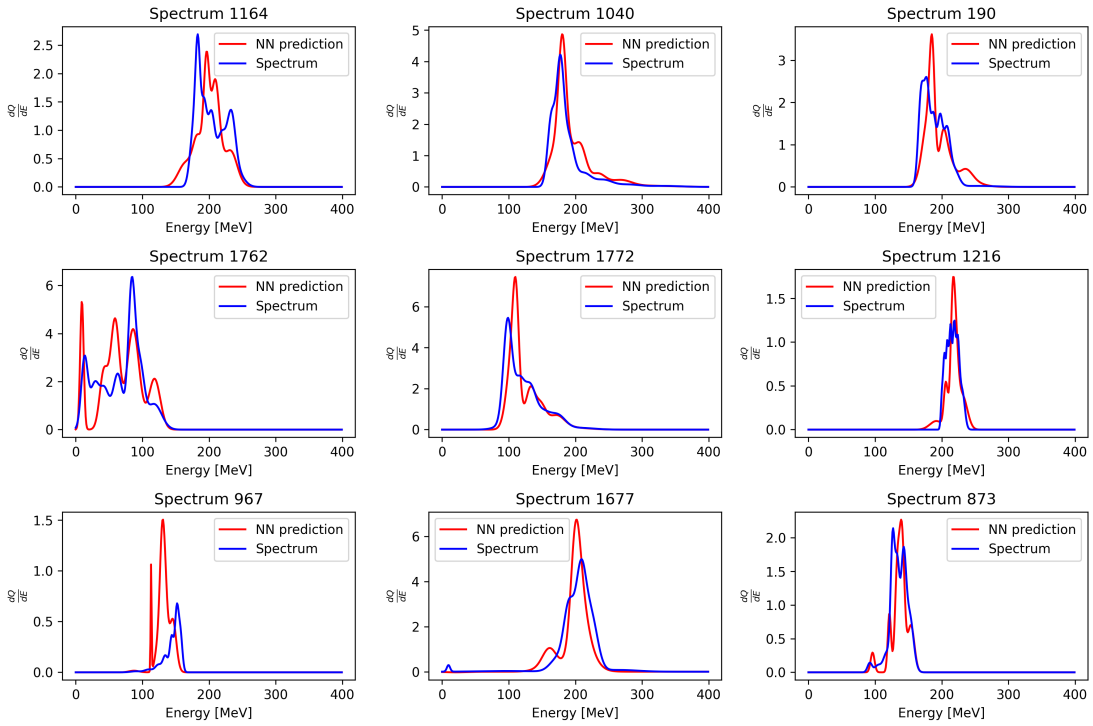

Figure 3.21: Nine randomly chosen spectra shown in blue and their corresponding neural network predictions in red using the MLP surrogate model.

A complementary approach, developed by my colleague Mykyta Lenivenko, focuses on the inverse problem: inferring input parameters from measured beam observables. This alternative method is not presented here, but it represents another promising avenue within the PALLAS team.

3.8 Conclusion and perspective

In conclusion, our in-depth numerical study focused on applying machine learning in the context of a laser plasma injector optimisation design. It demonstrates that a surrogate model approach is relevant for beam optimisation and stability, increasing efficiency and requiring a lower simulation cost. We successfully constructed models that exhibit high performance in predicting electron beam parameters.

We emphasised the importance of data distribution in achieving accurate results with SM. Our analysis has shown that R^2 scores converge rapidly towards 1 if trained on sufficiently uniform data sets. This capability enabled us to use the surrogate models effectively to identify optimal working points for designing LPI electron sources.

Furthermore, SM provides a comprehensive view of potential LPI beam parameters, facilitating the identification of stable operational regions and can drive the development of plasma targets for higher repetition rate laser-plasma accelerators with limited laser intensity. These models are straightforward to implement and can be continuously refined by incorporating new simulation data.

However, we identified certain limitations of the surrogate models. They tend to underperform in regions where data points are sparse and exhibit poorer performance at the lower end of the output range (low charge, low emittance and low energy spread). A potential improvement could be to add new simulation data in those regions and train our model on a larger interval of the output space than the test region.

Our study highlighted the ability of MLP and GP to generalise well and achieve the highest predictive performance among the ML methods considered. The LPI surrogate model can be used as an electron source for start-to-end simulation studies, opening the way to model a full accelerator beamline with variation in LPI electron source input parameters.

These promising results show that these methods could eventually be used with experimental data of the LPI or a hybrid version between experimental and simulation data since the time necessary to gather a large amount of experimental data is much shorter than PIC simulations. Such models could be implemented using a multi-fidelity approach as in [26]. More weight would be added to experimental data in comparison to simulations.

The next step would be to apply the SM within a control framework to introduce an efficient feedback loop for LPI beam stabilisation and control. This could be achieved either by using a reverse SM [27, 28] to map the output space back to the input space—although this task is numerically challenging since the four output parameters are not independent, and the existence and uniqueness of the solution are not guaranteed—or by embedding the forward SM into a model predictive control scheme [29] or a model-based reinforcement learning framework [30], both of which can directly exploit the predictive power of the SM to optimise control actions.

3.8.1 Output electron beam parameters evaluation

The electron beam features are retrieved using a postprocessing script based on APtools python code [15] for the last time step corresponding to the end of the electron plasma density longitudinal profile. The figure Fig. 3.22 shows an example of median energy E_{med}, median absolute deviation E_{mad} and charge corresponding to the integration of the electron beam energy distribution from the minimum energy tracked in the PIC code to the maximum energy.

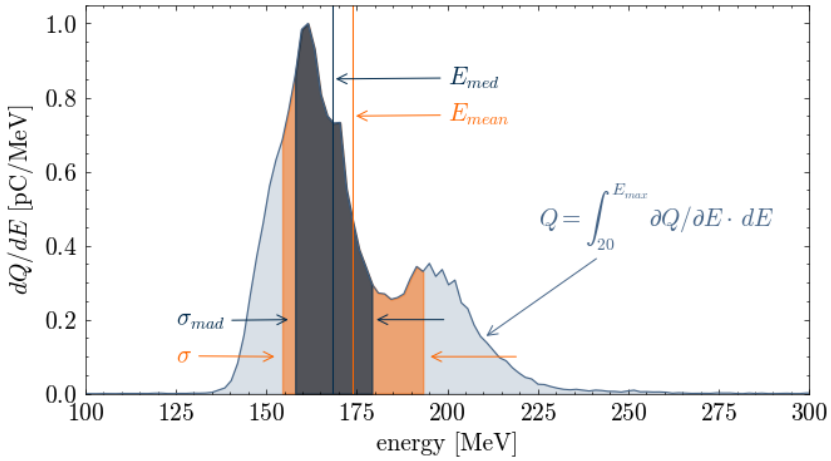

Figure 3.22: electron beam statistical features processed from the energy distribution

Only the trackParticles SMILEI openPMD[31] output is used to post-process electron beam features. The method used in the present paper can be extended to any PIC code data using openPMD format.

References

[1] P. Drobniak et al. “Random scan optimization of a laser-plasma electron injector based on fast particle-in-cell simulations”. In: *Phys. Rev. Accel. Beams* 26 (9 Sept. 2023), p. 091302. doi: 10.1103/PhysRevAccelBeams.26.091302. url: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.26.091302.

[2] Pierre Drobniak. “Development of a 150 MeV laser-plasma injector prototype”. Accelerator Physics [physics.acc-ph], English. NNT: 2023UPASP184, tel-04472687. PhD thesis. Université Paris-Saclay, 2023.

[3] J. Derouillat et al. “Smilei: A collaborative, open-source, multi-purpose particle-in-cell code for plasma simulation”. In: *Computer Physics Communications* 222 (2018), pp. 351–373.

[4] *Genci center for high performance computing*. https://www.genci.fr/en. Accessed: 2010-09-30.

[5] C.K. Birdsall and A.B. Langdon. *Plasma physics via computer simulation*. CRC press, 2018.

[6] A. Döpp et al. “Data-driven science and machine learning methods in laser–plasma physics”. In: *High Power Laser Science and Engineering* 11 (2023), e55.

[7] G. Carleo et al. “Machine learning and the physical sciences”. In: *Reviews of Modern Physics* 91.4 (2019), p. 045002.

[8] M. Leshno et al. “Multilayer Feedforward Networks with a Nonpolynomial Activation Function Can Approximate Any Function”. In: *Neural Networks* 6 (1993), pp. 861–867.

[9] N.Qian. “On the momentum term in gradient descent learning algorithms”. In: *Neural networks* 12.1 (1999), pp. 145–151.

[10] J.L.Ba D.P.Kingma. “A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* 1412.6 (2014).

[11] B. Shahriari et al. “Taking the Human Out of the Loop: A Review of Bayesian Optimization”. In: *Proceedings of the IEEE* 104.1 (2016), pp. 148–175. doi: 10.1109/JPROC.2015.2494218.

[12] M. Chen et al. “Theory of ionization-induced trapping in laser-plasma accelerators”. In: *Physics of Plasmas* 19.3 (2012).

[13] G. Golovin et al. “Tunable monoenergetic electron beams from independently controllable laser-wakefield acceleration and injection”. In: *Physical Review Special Topics-Accelerators and Beams* 18.1 (2015), p. 011301.

[14] P. Drobniak et al. *Two-chamber gas target for laser-plasma accelerator electron source*. 2023. arXiv: 2309.11921 [physics.acc-ph]. url: https://arxiv.org/abs/2309.11921.

[15] *APtools: Accelerator Physics tools in Python*. https://github.com/AngelFP/APtools. 2024.

[16] W. Lu et al. “Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime”. In: *Physical Review Accelerators and Beams* 10 (2007), p. 061301.

[17] M. Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. url: http://tensorflow.org/.

[18] K. He et al. “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification”. In: *Proceedings of the IEEE international conference on computer vision*. 2015, pp. 1026–1034.

[19] P. Refaeilzadeh, L. Tang, and H. Liu. “Cross-validation”. In: *Encyclopedia of database systems* (2009), pp. 532–538.

[20] T. Chen and C. Guestrin. “Xgboost: A scalable tree boosting system”. In: *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 2016, pp. 785–794.

[21] F. Pedregosa et al. “Scikit-learn: Machine learning in Python”. In: *the Journal of machine Learning research* 12 (2011), pp. 2825–2830.

[22] *LPA PIC simulations Data: LPI surrogate models*. https://gitlab.in2p3.fr/pallas/simulations/lpism. 2024.

[23] S. Jalas et al. “Bayesian optimization of a laser-plasma accelerator”. In: *Physical review letters* 126.10 (2021), p. 104801.

[24] F. Irshad et al. “Pareto Optimization of a Laser Wakefield Accelerator”. In: *Physical review letters* 133 (2024), p. 085001.

[25] S. Jalas et al. “Tuning curves for a laser-plasma accelerator”. In: *Physical Review Accelerators and Beams* 26.7 (2023), p. 071302.

[26] F. Irshad, S. Karsch, and A. Döpp. “Multi-objective and multi-fidelity Bayesian optimization of laser-plasma acceleration”. In: *Physical Review Research* 5.1 (2023), p. 013063.

[27] F. Bethke et al. *Invertible Surrogate Models: Joint surrogate modelling and reconstruction of Laser-Wakefield Acceleration by invertible neural networks*. arXiv:2106.00432 [physics]. June 2021. url: http://arxiv.org/abs/2106.00432 (visited on 04/19/2023).

[28] T. Miethlinger, N. Hoffmann, and T. Kluge. "Acceptance Rates of Invertible Neural Networks on Electron Spectra from Near-Critical Laser-Plasmas: A Comparison". In: *Parallel Processing and Applied Mathematics*. Ed. by R. Wyrzykowski et al. Cham: Springer International Publishing, 2023, pp. 273–284.

[29] M. Schwenzer et al. "Review on model predictive control: An engineering perspective". In: *The International Journal of Advanced Manufacturing Technology* 117.5 (2021), pp. 1327–1349.

[30] T. M. Moerland et al. "Model-based reinforcement learning: A survey". In: *Foundations and Trends® in Machine Learning* 16.1 (2023), pp. 1–118.

[31] *Open Standard for Particle-Mesh Data*. https://github.com/openPMD/openPMD-api. 2024.

4 -Spatio-Spectral Characterization and Phase Retrieval

4.1 Introduction

In this section, we present the calibration of a multispectral camera (MSCCD) and then use the calibrated data to perform phase retrieval on the simulated data. The calibration was performed at the Center for Advanced Laser Applications (CALA) in Munich.

4.2 Multi spectral camera

The multispectral camera (MSCCD) used in the following was manufactured by Spectral Devices Inc. It consists of a mosaic of 8 filters, each selecting a spectral band with the following wavelengths: [785, 793, 803, 812, 824, 832, 842, 853] nm.

The camera is equipped with a CCD chip of 2048×2048 pixels, with a pixel size of 5.5 µm. Each macropixel is composed of 8×8 pixels, resulting in a total of 256×256 macropixels on the chip. All 8 spectral filters are present in each macropixel, and their layout is shown in Fig. 4.1.

(1, 1)	(1, 2)	(1, 3)	(1, 4)	(1, 5)	(1, 6)	(1, 7)	(1, 8)
(2, 1)	(2, 2)	(2, 3)	(2, 4)	(2, 5)	(2, 6)	(2, 7)	(2, 8)
(3, 1)	(3, 2)	(3, 3)	(3, 4)	(3, 5)	(3, 6)	(3, 7)	(3, 8)
(4, 1)	(4, 2)	(4, 3)	(4, 4)	(4, 5)	(4, 6)	(4, 7)	(4, 8)
(5, 1)	(5, 2)	(5, 3)	(5, 4)	(5, 5)	(5, 6)	(5, 7)	(5, 8)
(6, 1)	(6, 2)	(6, 3)	(6, 4)	(6, 5)	(6, 6)	(6, 7)	(6, 8)
(7, 1)	(7, 2)	(7, 3)	(7, 4)	(7, 5)	(7, 6)	(7, 7)	(7, 8)
(8, 1)	(8, 2)	(8, 3)	(8, 4)	(8, 5)	(8, 6)	(8, 7)	(8, 8)

Figure 4.1: Disposition of the filters on a macropixel of the multispectral camera. Each color represents a different filter

4.3 Calibration

We can model the signal recorded by the camera as follows:

$$MS_{\lambda_k}(x, y) = \int \beta_{\lambda_k}(x, y, \lambda) I(x, y, \lambda) d\lambda \tag{4.1}$$

where $I(x, y, \lambda)$ is the spatio-spectral intensity, $\beta_{\lambda_k}(x, y, \lambda)$ is the spatio-spectral response of the MSCCD, and λ_k is the wavelength associated with each channel.

Due to manufacturing imperfections and pre-calibration inaccuracies, each measurement is affected by a spatio-spectral gain. The goal of the calibration is to apply a correction for each macro-pixel such that $\int \beta_{\lambda_k}(x, y, \lambda)d\lambda = 1 \quad \forall(x, y)$. Furthermore, since the bandwidth of each spectral filter is not infinitesimally narrow, there is residual cross-talk between adjacent channels. The calibration also aims to reduce this cross-talk.

4.3.1 Calibration setup

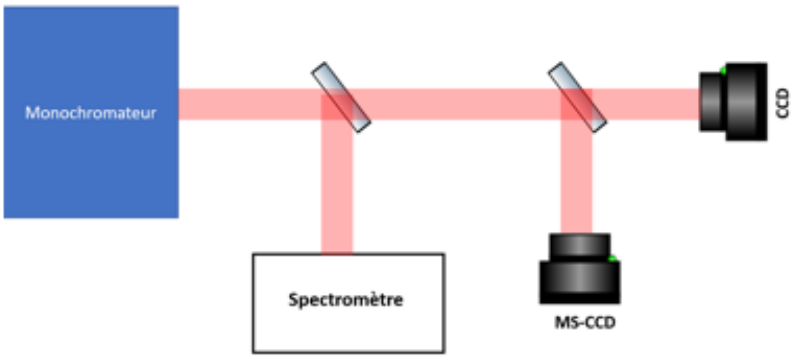

Figure 4.2: Scheme of the calibration setup

In Figure (4.2), a scheme of the calibration setup is presented. The monochromator delivers a quasi-monochromatic beam with a full-width at half-maximum bandwidth of 2.5 nm. In total, we made 220 acquisitions with beams of wavelength varying from 770 nm to 879.5 nm. The beam is then sent to the spectrometer in order to know the intensity $\mathcal{I}(\lambda)$ delivered. Finally, the beam is sent to the MSCCD and the classical CCD.

4.3.2 Image processing

In an ideal calibration setup, the incoming intensity on the MSCCD would be uniform. In this ideal case, the classical CCD would be unnecessary since $I(\lambda)$ is known from the spectrometer, and we would have direct access to the spatio-spectral response of the MSCCD. However, in our setup, this was not the case, so we used a CCD to obtain information about $I(x, y, \lambda)$. The response of the CCD is given by:

$$CCD(x, y) = \int \alpha(\lambda) I(x, y, \lambda) d\lambda$$

Since the MSCCD and CCD are imaging the same plane, we can simulate what the MSCCD would see under uniform illumination by computing the proportion of photons striking the cameras. For each acquired image, the proportion of photons is:

$$p(x, y, \lambda_i) = \frac{\alpha(\lambda_i) I(x, y, \lambda_i)}{\sum_{x,y} \alpha(\lambda_i) I(x, y, \lambda_i)}$$

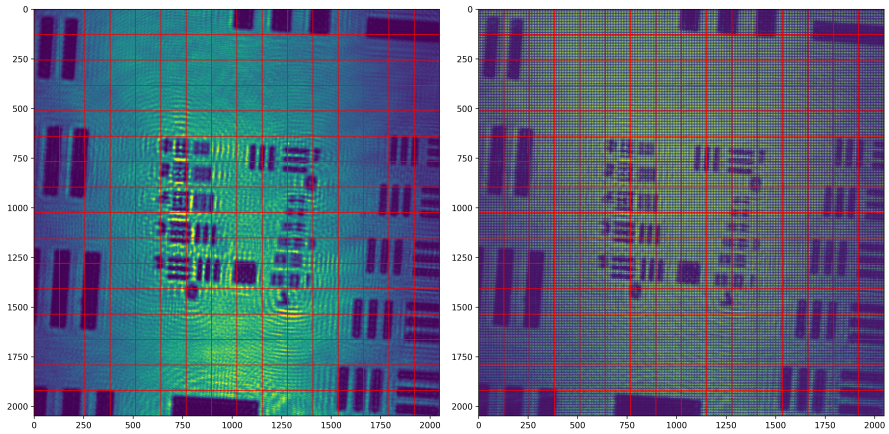

Figure 4.3: Image used for calibration. On the left is the CCD image; on the right, the MSCCD. The photon proportion is calculated within the red squares.

Here, i refers to the i-th acquired image, which corresponds to the i-th narrowband illumination sent by the monochromator, centered at wavelength λ_i. Each image is taken under a different λ_i, allowing us to sample the system's response across multiple discrete wavelengths (the integral over λ is omitted since the incoming beam is considered monochromatic). In practice, $p(x, y, \lambda_i)$ is not computed at every macropixel position (x, y); instead, the proportion is assumed to be constant over groups of 16×16 macropixels. This simplification accounts for the fact that the CCD and MSCCD images do not perfectly overlap due to slight differences in alignment and imaging geometry. Additionally, it reduces the impact of high-frequency noise caused by sensor noise, local artefacts, or small optical mismatches, leading to a more robust estimation of photon distribution across the sensors.

The intensity acquired by the MSCCD under illumination at wavelength λ_i is first corrected for spatial inhomogeneity by rescaling with the inverse of the photon proportion $p(x, y, \lambda_i)$ estimated from the CCD image. This yields an intermediate image:

$$MS2_{\lambda_k}(x, y, \lambda_i) = \frac{MS_{\lambda_k}(x, y, \lambda_i)}{p(x, y, \lambda_i)} = \beta_{\lambda_k}(x, y, \lambda_i) \cdot I(\lambda_i)$$

Here, $MS_{\lambda_k}(x, y, \lambda_i)$ is the raw signal measured by the MSCCD at spectral band λ_k when illuminated at λ_i, and $I(\lambda_i)$ is the total spectral intensity, assumed spatially uniform after correction by $p(x, y, \lambda_i)$.

To simulate the response of the MSCCD under uniform illumination, we then normalise $MS2_{\lambda_k}$ by the spectrally-resolved intensity $\mathcal{I}(\lambda_i)$ measured independently by the spectrometer:

$$MS2_{\lambda_k}(x, y, \lambda_i) \xrightarrow{\div \mathcal{I}(\lambda_i)} \beta_{\lambda_k}(x, y, \lambda_i)$$

This step assumes that the CCD-estimated intensity distribution $p(x, y, \lambda_i)$ correctly accounts for spatial variations, so that $I(\lambda_i)$ and the spectrometer-measured $\mathcal{I}(\lambda_i)$ are proportional up to a constant factor. This allows us to isolate the intrinsic response $\beta_{\lambda_k}(x, y, \lambda_i)$ of the MSCCD independently of illumination power.

The image corrected can then be calibrated using a method similar to that of Dorrer *et al.* [1]. Since the multi-spectral camera has 8 spectral channels, we use measurements taken at the 8 corresponding wavelengths to construct a calibration matrix. This results in an 8×8 matrix, where each element is defined as:

$$P_{k,j} = \beta_{\lambda_k}(x, y, \lambda_j)$$

Here, λ_k and λ_j refer to the central wavelengths of the 8 spectral filters of the MSCCD.

The measured signal at each pixel (x, y) across all spectral channels can then be expressed as:

$$\begin{bmatrix} MS2_{\lambda_1}(x, y, \lambda_j) \\ MS2_{\lambda_2}(x, y, \lambda_j) \\ \vdots \\ MS2_{\lambda_8}(x, y, \lambda_j) \end{bmatrix} = \begin{bmatrix} P_{1,1} & P_{1,2} & \cdots & P_{1,8} \\ P_{2,1} & P_{2,2} & \cdots & P_{2,8} \\ \vdots & \vdots & \ddots & \vdots \\ P_{8,1} & P_{8,2} & \cdots & P_{8,8} \end{bmatrix} \begin{bmatrix} \delta_{1j} \\ \delta_{2j} \\ \vdots \\ \delta_{8j} \end{bmatrix} \tag{4.2}$$

Here, δ_{kj} is the Kronecker delta, equal to 1 if $k = j$ and 0 otherwise. This formulation isolates the calibration response for a specific illumination wavelength, λ_j.

For each macropixel (x, y), we compute the inverse of the calibration matrix, whose elements are denoted by $P^{-1}_{k,j}$. Applying this inverse matrix to the measured signal allows us to suppress cross-talk between channels at the calibration points. Specifically, it ensures that if the incoming light is monochromatic with wavelength $\lambda = \lambda_k$, the corrected measurement will reflect the true intensity at that wavelength:

$$MS3_{\lambda_k}(x, y) = \sum_{j=1}^{8} P^{-1}_{k,j} \cdot MS2_{\lambda_j}(x, y)$$

The calibrated spectral response of a uniform illumination is presented in Fig. 4.4.

This calibration procedure assumes that the incoming light is either monochromatic (i.e., $\lambda = \lambda_k$) or a linear combination of delta functions centered on the calibration wavelengths λ_k. In practice, the incoming beam can have a continuous spectral distribution.

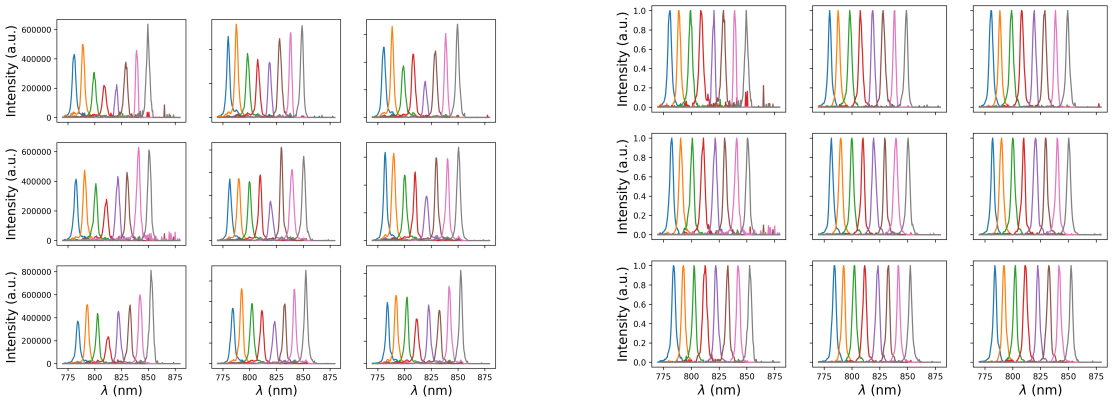

(a) $MS2_{\lambda_k}(x, y)$ before calibration

(b) $MS3_{\lambda_k}(x, y)$ after calibration

Figure 4.4: Spectral response of the camera on 9 different macropixels before and after calibration. The top-left plot corresponds to the top-left macropixel of the MSCCD, while the bottom-right plot corresponds to the bottom-right macropixel.

For instance, a flat input spectrum (e.g., a Heaviside function over the relevant bandwidth) should ideally yield a flat calibrated output. However, due to variations in the integrated spectral response across channels and sensor position, caused by differences in filter bandwidth and shape, a simple matrix inversion does not ensure a flat response.

To correct for this, we define the system's response to a flat spectral input as:

$$G_{\lambda_k}(x, y) = \int \sum_{j=1}^{8} P_{k,j}^{-1} \beta_{\lambda_j}(x, y, \lambda) d\lambda \tag{4.3}$$

This gain function corresponds to the reconstructed output at channel λ_k, assuming the input spectrum is flat across the relevant bandwidth. To ensure that the reconstructed signal is spectrally flat under such conditions, we normalize the reconstruction by this gain factor. The final calibrated image is therefore given by:

$$MScalib_{\lambda_k}(x, y) = \frac{\sum_{j=1}^{8} P_{k,j}^{-1} \cdot MS2_{\lambda_j}(x, y)}{G_{\lambda_k}(x, y)} \tag{4.4}$$

This ensures that if the input spectrum is flat, the reconstructed signal remains flat across all channels, as desired.

4.4 Multispectral camera sensitivity

One of the main advantages of the multispectral camera (MSCCD) is that it provides a highly practical tool for aligning the laser pulse compressor. In particular, it enables

real-time visualization of any residual angular dispersion and allows one to distinguish between chromatic aberrations and purely spatial aberrations of the laser beam.

By combining the MSCCD with the FROG diagnostic, it is possible to achieve precise adjustment of the optical grating compressor at 10 Hz. To evaluate the sensitivity of our setup, we intentionally misaligned the compressor over a range from -30 mdeg to $+30$ mdeg. The corresponding measurements are presented in Fig. 4.6.

First, we developed a graphical user interface that allows real-time visualization of the different spectral channels. A screenshot of this interface is shown in Fig. 4.5. With our setup, using a $\times 50$ microscope objective, we achieve a sensitivity of $6\,\mu\mathrm{m\,mdeg}^{-1}$.

Figure 4.6 shows the multispectral camera response for different compressor misalignment angles. The images are displayed in RGB format for visualization purposes. The top row presents the RGB representation, where each color corresponds to a different spectral channel of the MSCCD. The bottom row shows the summed intensity over all channels.

Finally, Fig. 4.7 compares the sensitivity of a conventional monochrome CCD camera (blue curve) with that of the multispectral camera (orange curve). The orange curve represents the sum of the centroid displacements measured across all spectral channels of the MSCCD, whereas the blue curve corresponds to the beam waist variation measured with a standard CCD camera with a magnification of 2.8. The main advantage of this setup in contrast with the traditional camera is that it allows us to decipher between spatial and chromatic aberrations.

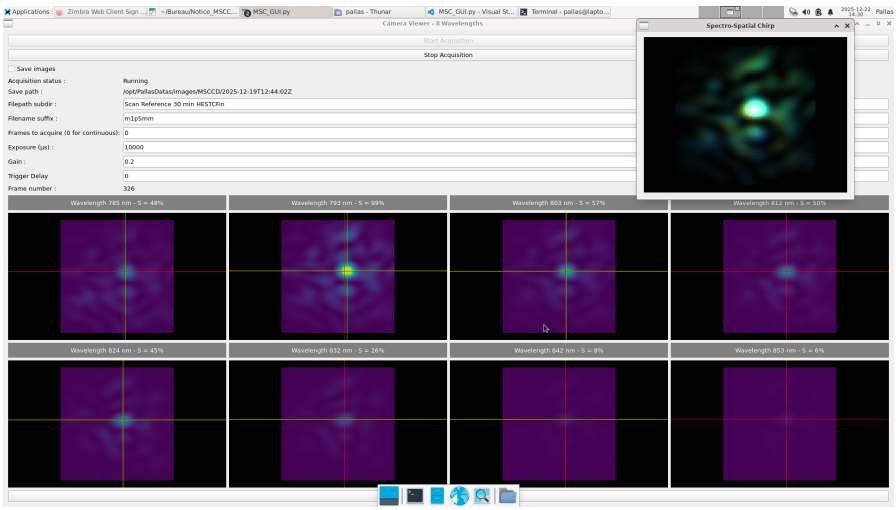

Figure 4.5: Screenshot of the graphical user interface with the compressor misaligned by $+5$ mdeg. The red cross indicates the centroid position of the $812\,\mathrm{nm}$ channel, while the green crosses correspond to the centroid positions of the other spectral channels. The "spectro-spatial chirp" panel represents the summed intensity over all channels.

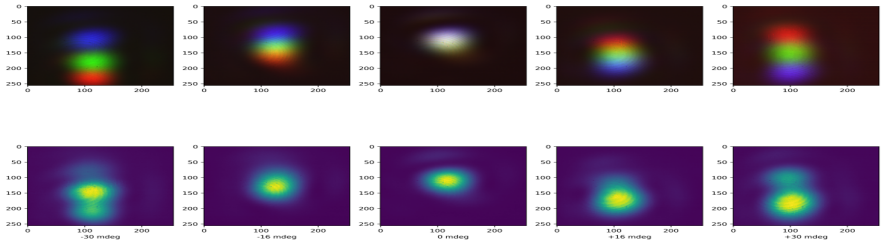

Figure 4.6: Multispectral camera measurements for different compressor misalignment angles. Top row: RGB representation, where each color corresponds to a different spectral channel of the MSCCD. Bottom row: summed intensity over all channels.

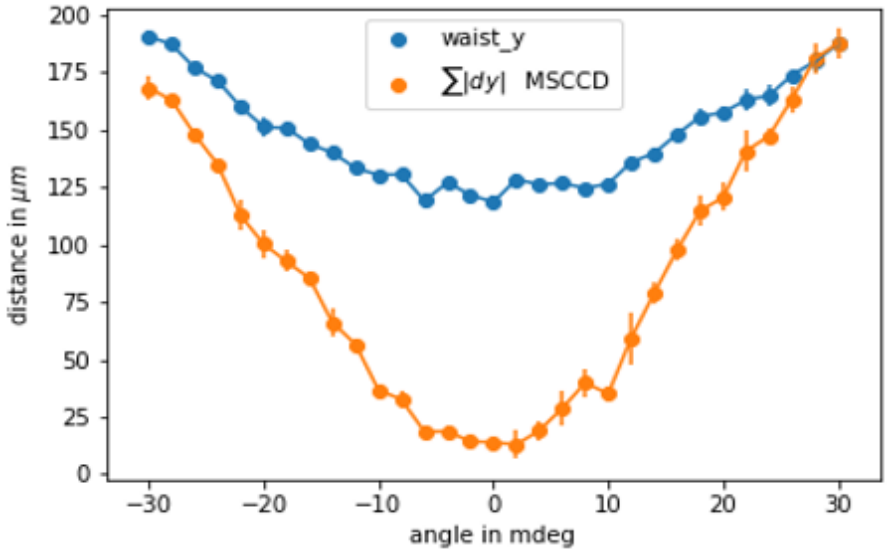

Figure 4.7: Comparison of the sensitivity of a conventional monochrome CCD camera (blue) and the multispectral camera (orange). The orange curve represents the summed centroid displacement measured with the MSCCD, while the blue curve corresponds to the beam waist variation measured with the CCD camera.

4.5 Phase retrieval

Before discussing phase retrieval, we first present one propagator and the phase retrieval algorithm that will be used throughout the rest of the thesis.

4.5.1 Angular spectrum decomposition

In scalar diffraction theory, the exact solution to the Helmholtz equation is the Rayleigh–Sommerfeld integral [2], which is derived using a Green-function approach. An equivalent solution [2] is obtained by considering the angular spectrum of the electric field. By angular spectrum, we mean:

$$E(k_x, k_y) = \iint_{-\infty}^{+\infty} E(x, y, 0)\, \exp[-\mathrm{i}\,(k_x x + k_y y)]\, \mathrm{d}x\, \mathrm{d}y. \tag{4.5}$$

The above integral represents the spatial frequency content of the electric field at the position $z = 0$, where each component is a plane wave propagating at a different angle. After propagation over a distance z, each component of the angular spectrum accumulates a different phase, denoted by:

$$H(k_x, k_y) = \exp\left(\mathrm{i}k_z z\right), \tag{4.6}$$

with

$$k_z = \sqrt{k^2 - k_x^2 - k_y^2}, \quad k = \frac{\omega}{c}. \tag{4.7}$$

After propagation over a distance z, the electric field can be written as:

$$E(x, y, z) = \iint_{-\infty}^{+\infty} E(k_x, k_y, 0)\, H(k_x, k_y)\, \exp[\mathrm{i}\,(k_x x + k_y y)]\, \mathrm{d}k_x\, \mathrm{d}k_y. \tag{4.8}$$

An important caveat is that if $k_x^2 + k_y^2 > k^2$, the wave component becomes evanescent and carries no energy. For our purposes, we only consider components where $k_x^2 + k_y^2 < k^2$.

The second important consideration is the numerical implementation of this method. Unlike a simple Fast Fourier Transform, the angular spectrum method has two Nyquist limits: the propagator itself is subject to a sampling constraint. To avoid aliasing, following the prescription by Kyoji Matsushima et al.[3], the maximum transverse spatial frequencies along x and y are given by:

$$k_{\max,x} = \frac{2\pi}{\lambda\sqrt{\left(\frac{\Delta k_x\,|z|}{\pi}\right)^2 + 1}}, \quad k_{\max,y} = \frac{2\pi}{\lambda\sqrt{\left(\frac{\Delta k_y\,|z|}{\pi}\right)^2 + 1}}, \tag{4.9}$$

where λ is the wavelength, Δk_x and Δk_y are the sampling intervals in the spatial frequency domain along x and y, and z is the propagation distance. These limits ensure that the angular spectrum method is numerically stable and avoids aliasing of high spatial frequency components. A simple way to increase the maximum transverse spatial frequency is to add zero-padding to the electric field before propagation. Unless otherwise specified, the propagation method used throughout the rest of this thesis is the angular spectrum decomposition.

4.5.2 Gerchberg–Saxton algorithm

The Gerchberg–Saxton Algorithm (GSA) is a phase retrieval method that relies on an alternating projection between multiple constraint sets. In optics, and in particular for the phase retrieval of laser beams, the algorithm requires at least two constraint sets defined in different transverse planes.

We consider two planes located at longitudinal positions z_1 and z_2. The associated constraint sets are defined as

$$A = \Big\{ E(x, y, z = z_1) \in \mathcal{H} \;\Big|\; |E| = |E|_{M,1} \Big\}, \tag{4.10}$$

$$B = \Big\{ E(x, y, z = z_2) \in \mathcal{H} \;\Big|\; |E| = |E|_{M,2} \Big\}, \tag{4.11}$$

where $|E|_{M,1}$ and $|E|_{M,2}$ denote the measured field amplitude in planes z_1 and z_2, respectively, and $\mathcal{H}$ is the Hilbert space of admissible complex-valued fields.

The projection onto each constraint set is defined as

$$P_A(F) = \arg\min_{E \in A} \big|E - F\big|, \quad P_B(F) = \arg\min_{E \in B} \big|E - F\big|, \tag{4.12}$$

where F denotes the current field estimate before applying the constraint, and the projection returns the closest field E that satisfies the corresponding constraint. In practice, for optical phase retrieval, these projections simply consist of replacing the modulus of the current field estimate by the measured modulus while preserving its phase.

Starting from an initial guess in plane z_1, $E_{z_1}^{(0)}$, the Gerchberg–Saxton algorithm iterates according to the following loop:

$$E_{z_2}^{(n)} = P_r\Big(P_A\left(E_{z_1}^{(n)}\right)\Big), \tag{4.13}$$

$$P_B\left(E_{z_2}^{(n)}\right) = |E|_{M,2}\, \exp\Big[i\, \arg\Big(E_{z_2}^{(n)}\Big)\Big], \tag{4.14}$$

$$E_{z_1}^{(n+1)} = P_r^{-1}\Big(P_B\left(E_{z_2}^{(n)}\right)\Big), \tag{4.15}$$

$$P_A\left(E_{z_1}^{(n+1)}\right) = |E|_{M,1}\, \exp\Big[i\, \arg\Big(E_{z_1}^{(n+1)}\Big)\Big]. \tag{4.16}$$

Here, P_r denotes the propagation between planes. For closed convex constraint sets, the convergence of the algorithm is guaranteed [4]. However, in our optical problem, the sets A and B are non-convex, and convergence is not guaranteed. Fortunately, since our fields are physical and, in most cases, well-behaved, the algorithm converges reliably in practice.

4.5.3 Phase retrieval on simulated data

Now that we have a calibrated signal, we can use equation (4.4) to generate the response of the MSCCD under an arbitrary illumination $I(x, y, \lambda)$. This simulated response can then be used to perform some phase retrieval using a Gerchberg-Saxton algorithm (GSA). The intensity patterns we will simulate are beams at and near focus with pulse-front tilt (PFT) induced by an input beam with angular dispersion (AD). The equation of the input beam is the following :

$$E(x', y', \lambda) = E_0 \exp\left(\frac{-(x'^2 + y'^2)}{w_0^2}\right) \exp\left(\frac{-(\lambda - \lambda_0)^2}{\Delta\lambda^2}\right) e^{ik\sin(\theta(\lambda))x'} \tag{4.17}$$

This beam is subsequently focused by a lens, and we can represent it using the Fresnel diffraction integral:

$$E(x, y, z) = \frac{e^{ikz}}{i\lambda z}\, e^{i\frac{k}{2z}(x^2+y^2)} \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} E(x', y')\, e^{i\frac{k}{2}(x'^2+y'^2)\left(\frac{1}{z}-\frac{1}{f}\right)}\, e^{-i\frac{k}{z}(xx'+yy')}\, dx'\, dy' \tag{4.18}$$

The solution yields:

$$\begin{aligned} E(x, y, z) = \frac{e^{ikz}}{i\lambda z}\, e^{i\frac{k}{2z}(x^2+y^2)} \exp\left(i\arctan\left(\frac{z_r(f-z)}{zf}\right)\right) \\ \times \frac{\pi w_0^2 \exp\left[\frac{-\left[\left(\frac{kx}{z} - k\sin(\theta(\lambda))\right)^2 + \left(\frac{ky}{z}\right)^2\right] w_0^2 \left(1 + i\frac{z_r(f-z)}{zf}\right)}{4\left(1 + \left(\frac{z_r(f-z)}{zf}\right)^2\right)}\right]}{\sqrt{1 + \left(\frac{z_r(f-z)}{zf}\right)^2}} \end{aligned} \tag{4.19}$$

The parameters used in the simulation are the following, $f = 1$m, $w_0 = 2.5$mm $\theta(\lambda) \in [-0.17, 0.32]$ mrad, $\Delta\lambda = 80$ nm,$\lambda_0 = 808$ nm. To perform the GSA, we select three planes: one 5 cm before the focus, one at the focus, and one 5 cm after the focus. In Figure 4.8, the spatial intensity distributions at the three planes are shown. The top panels clearly reveal the spatial chirp induced by angular dispersion. The bottom panels display the simulated intensity signals from the 785 nm and 852 nm channels, as detected by the MSCCD's simulated response. These intensities are then used as inputs to perform the GSA.

In Figure 4.9, the spatial phase retrieved for the 785 nm channel is presented. The top panels show the ground-truth phase, while the bottom panels show the phase retrieved via GSA. The retrieved patterns closely match the true distributions. As expected, the GSA reconstructs a phase curvature in the forward and backward planes and a nearly flat phase at focus. For the reconstructed GSA phase we impose the center of

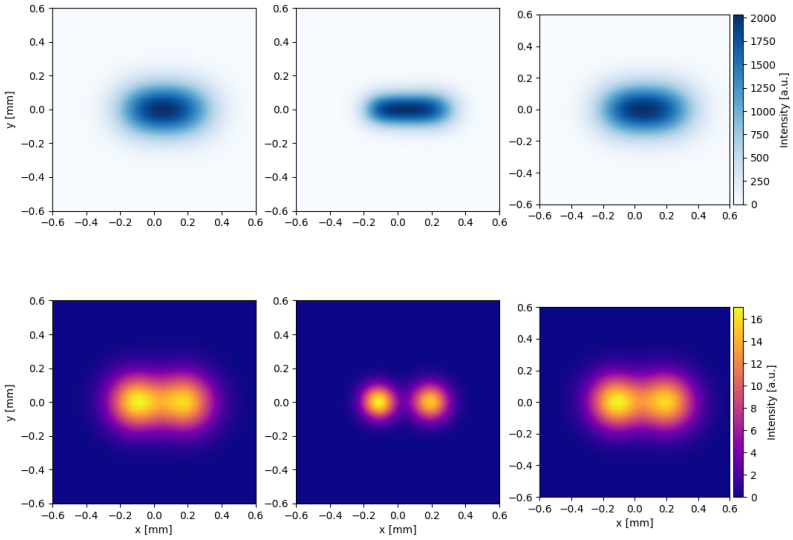

Figure 4.8: Intensity profiles of the AD/PFT beam, shown from left to right: backward, focus, and forward planes. The top panels present the theoretical solution, while the bottom panels show the MSCCD-detected intensities at 785nm and 852nm

the curvature phase to be equal t $-\pi/2$ for the backward plane and $\pi/2$ for the forward plane.

The final figure (where we neglected the Fresnel phase prefactor) displays the spatio-spectral phase, defined as:

$$\phi(x, \lambda) = \phi(x, y = 0, \lambda)$$

Again, the top panels show the true phase, while the bottom panels display the GSA-retrieved phase. Due to camera limitations, the spectral resolution in the bottom panels is limited to eight discrete wavelengths. The reconstruction is again really close to the real phase. The black line indicates the theoretical spatial chirp slope in units of $\mu\mathrm{m}/\mathrm{nm}$:

$$\alpha = \frac{zd}{f\,\delta\lambda},$$

where d is the observed spatial chirp measured at a wavelength offset $\delta\lambda$ from the reference wavelength. Despite small imperfections in the reconstructed phase, relevant information on angular dispersion (AD) and pulse-front tilt (PFT) can still be extracted from the observed slope.

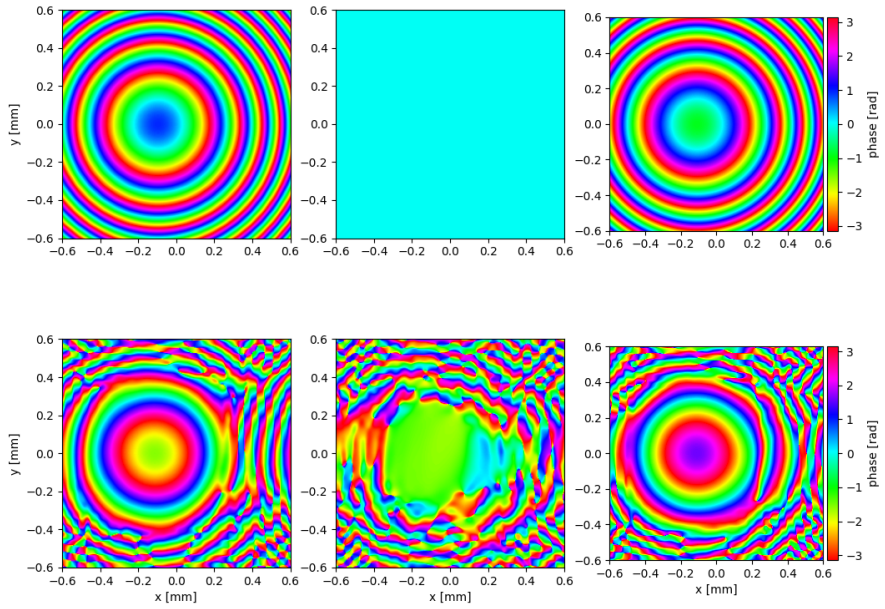

Figure 4.9: Phase profiles of the AD/PFT beam shown from left to right: backward, focus, and forward planes. The top panels are the theoretical solution, the bottom panels show the phase retrieved by the GSA at the MSCCD for 785 nm

Moreover, in this simple case, the reconstructed phase provides additional information compared to simply tracking the centroid position at focus. In particular, it allows one to discriminate between a beam that exhibits purely spatial chirp (SC) and a beam that exhibits purely angular dispersion (AD), which corresponds to the case analyzed here.

By a beam with pure spatial chirp, we mean a field of the form

$$E(x', y', \lambda) = E_0 \exp\left(-\frac{(x' - x_c(\lambda))^2 + y'^2}{w_0^2}\right) \exp\left(-\frac{(\lambda - \lambda_0)^2}{\Delta\lambda^2}\right), \tag{4.20}$$

where $x_c(\lambda)$ denotes the wavelength-dependent centroid position.

After propagation, the spatial phase $\phi(x, \lambda)$ appearing in the exponential for such a

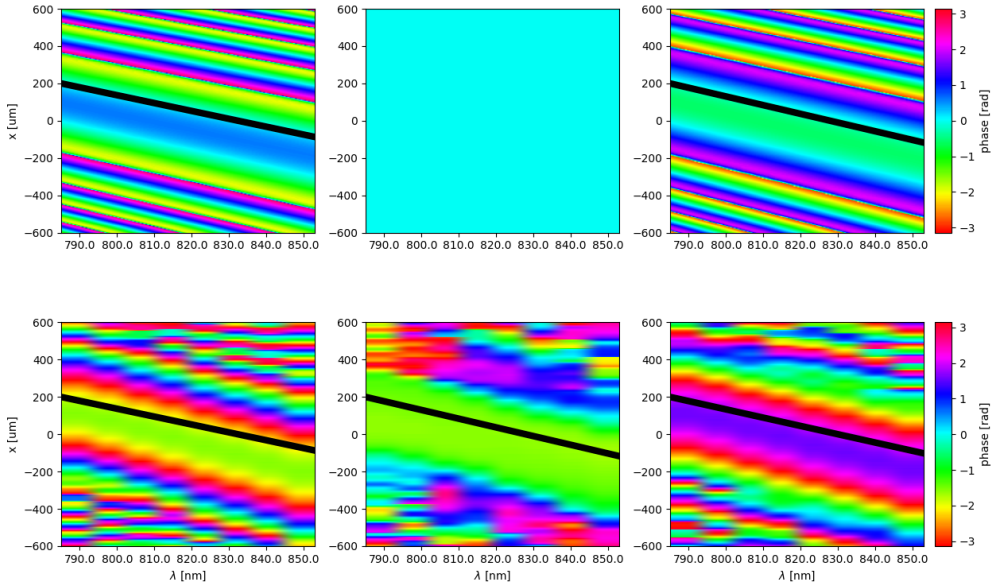

Figure 4.10: Spatio-spectral-Phase profiles of the AD/PFT beam shown from left to right: backward, focus, and forward planes. The top panels are the theoretical solution, the bottom panels show the phase retrieved by the GSA at the MSCCD. The black line indicates the theoretical spatial chirp slope.

beam can be written as

$$\phi(x,\lambda) = \left[\frac{-\left(\frac{k^2x^2}{z^2} - \frac{4x_c(\lambda)^2}{w_0^4}\right) w_0^2 \left(\frac{z_r(f-z)}{zf}\right)}{4\left(1+\left(\frac{z_r(f-z)}{zf}\right)^2\right)} - \frac{2kxx_c(\lambda)/z}{4\left(1+\left(\frac{z_r(f-z)}{zf}\right)^2\right)} \right]. \tag{4.21}$$

This expression shows that a purely spatially chirped beam gives rise to a phase containing both quadratic and linear terms in x. In contrast, in the case of pure angular dispersion, the phase remains purely quadratic, which allows experimental discrimination.

4.6 Discussion

The MSCCD demonstrates interesting performance in tracking laser intensity distributions across different wavelengths. Its spectral response shows minimal crosstalk between channels and, after calibration, provides high-fidelity relative-intensity measurements, as illustrated in Figure 4.4. However, several challenges must be addressed before the system can be reliably used in real-time experimental settings.

Firstly, the spectral bandwidth used in the simulation was broader than that available in the experimental setup. In fact, with our 32 ± 1 nm FWHM laser spectrum, only six channels are effectively usable in practice. The second issue is the spatial non-uniformity of the spectral filters across the sensor. As shown in Figure 4.11, the peak transmission wavelength for each channel varies depending on the position on the sensor. Although the mean peak wavelengths align with the specifications provided by Spectral Devices Inc., these local variations can introduce systematic errors in measurements.

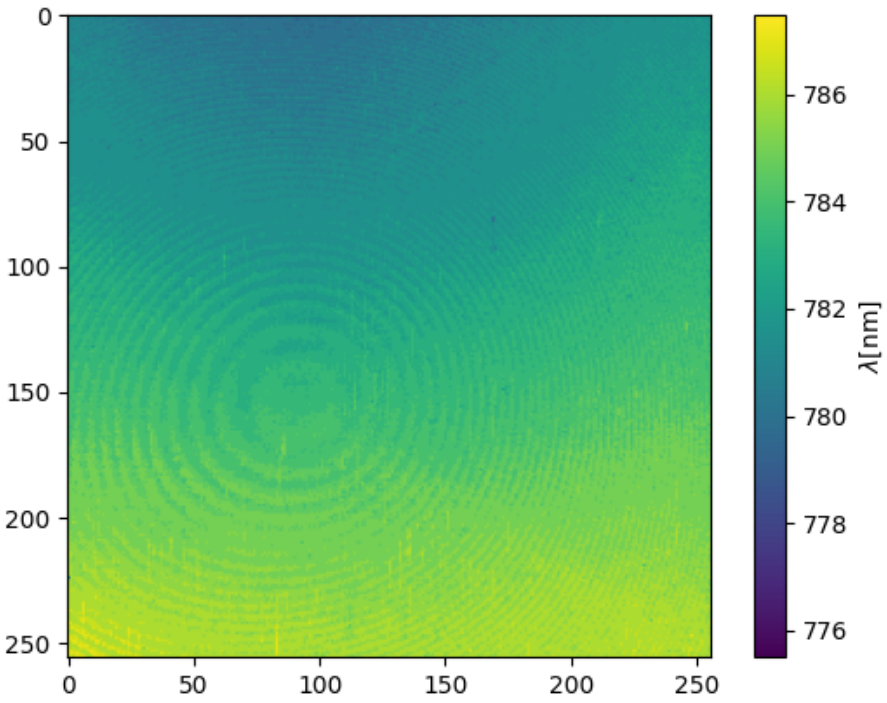

Figure 4.11: Distribution of the peak transmission wavelength of the multispectral camera at the specified value of 785 nm

GSA is relatively simple to implement and computationally efficient, and it proves to be effective in reconstructing simple phase profiles, such as the parabolic phase observed in the AD/PFT beam. However, because many spatio-temporal aberrations are encoded in the centroid positions of the intensity distributions, phase reconstruction is highly sensitive to beam pointing jitter. This sensitivity makes the accuracy of GSA strongly dependent on the stability and precision of the experimental alignment, which must be carefully controlled or independently characterized. For further improvement, particularly to enhance spectral resolution beyond the discrete channels of the MSCCD, physics-informed neural networks offer a promising approach for algorithmic interpolation or refinement. [5, 6]

4.7 Algorithmic development for phase retrieval

4.7.1 Dataset generation

As discussed earlier, the GSA can be highly effective for phase reconstruction when the optical setup is properly aligned. However, it is not suitable for real-time diagnostics, as the iterative procedure typically takes several seconds to minutes to converge. This becomes even more problematic when the algorithm must be applied simultaneously to the eight different channels of our system.

To address this limitation, we propose a Physics-Informed Neural Network (PINN) architecture that offers two major advantages over the GSA. First, the inference time is significantly faster, allowing real-time or near–real-time phase retrieval. Additionally, the PINN has the advantage of natively predicting the unwrapped phase, which is highly beneficial for feedback-control schemes. Since unwrapping errors can introduce artifacts, especially for complex wavefronts.

We thus propose a physics-informed neural network (PINN) to perform this task, inspired by the work in holography [7, 8] Since the multispectral camera was not operational for training such a PINN, we first generated a large amount of numerical data. The data were produced using the Fresnel diffraction integral, Eq. (4.18). The input near field

$$E(x', y')$$

was chosen to be a super-Gaussian beam of order 2, with wavelength $\lambda = 810$ nm and waist $w_{\mathrm{NF}} = 12.9$ mm, selected such that the non-aberrated beam has a waist of $20\,\mu$m after a focal length of 1 m.

We subsequently added aberrations sampled from the Shack-Hartmann (SH) sensor of the diagnostic table. To do so, we extracted 15 phase maps of the SH with RMS $\sigma < 0.3\,\mu$m (see Fig. 4.12). These maps correspond to the residual wavefront error after correction by the deformable mirror.

This set of 15 wavefronts forms the basis for the aberrations used to train the PINN, as they represent the residual aberrations remaining after wavefront correction. To introduce statistical variability into the training data, we first decomposed each wavefront into the first 36 Zernike modes. From these decompositions, we extracted the mean and standard deviation of each coefficient across the 15 measured maps.

New synthetic phase maps were then generated by sampling 36 Zernike coefficients from the corresponding distributions shown in Fig. 4.13. These sampled coefficients were used to reconstruct phase maps that serve as aberration inputs for the PINN training dataset.

Using this coefficient distribution, we generated 5000 synthetic phase maps. Figure 4.14 presents six randomly selected examples from this dataset.

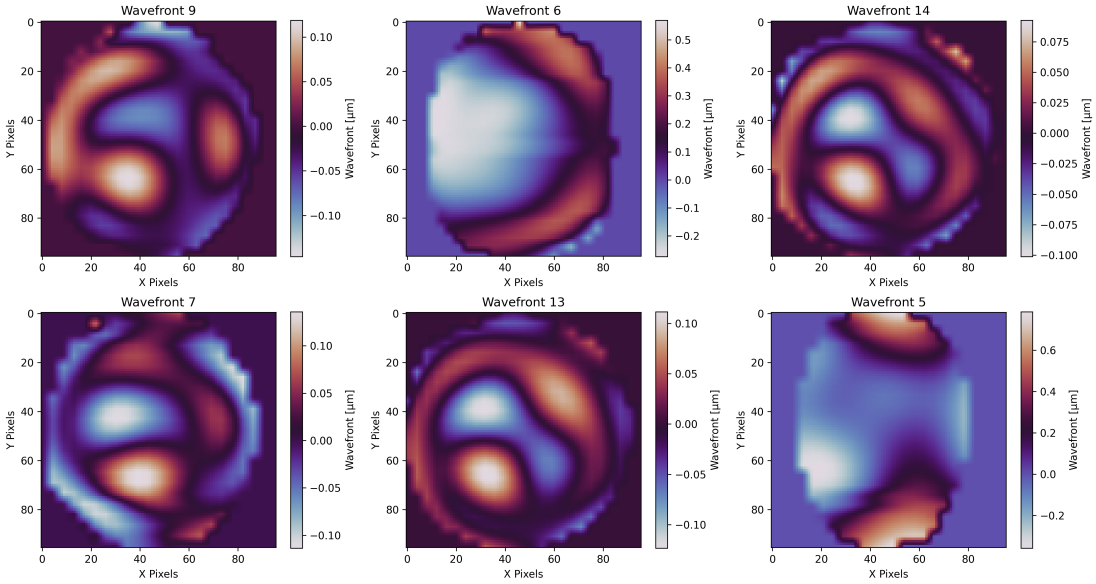

Figure 4.12: six randomly sampled wavefronts from the 15 chosen with rms error $\sigma < 0.3\mu m$

Finally, using Eq. (4.18), we propagated the near-field beam

$$E(x', y') = \exp\left[-\left(\frac{x^2 + y^2}{w_{\mathrm{NF}}^2}\right)^2\right] e^{i\phi_{\mathrm{input}}},$$

to three axial planes: one located 2 mm before the focus, one at the focus, and one 2 mm after the focus. In this configuration, the Rayleigh length is $z_r = 1.3mm$. Here, ϕ_{input} denotes one of the 5000 phase profiles drawn from Fig. 4.13 and illustrated in Fig. 4.14.

4.7.2 Machine learning architecture

We adopted the same architecture as presented in [8], which addressed a similar phase retrieval problem and demonstrated strong performance with their deep neural network. The chosen architecture is a U-Net with skip connections and residual blocks (see Figure 4.15 for a schematic representation).

The network is structured as follows:

- **Inputs:** The network receives three amplitude planes at and near the focus as input.
- **Encoder:** The encoder consists of five layers. Each layer applies a 3×3 2D convolution followed by a ReLU activation. The output is then passed through a residual

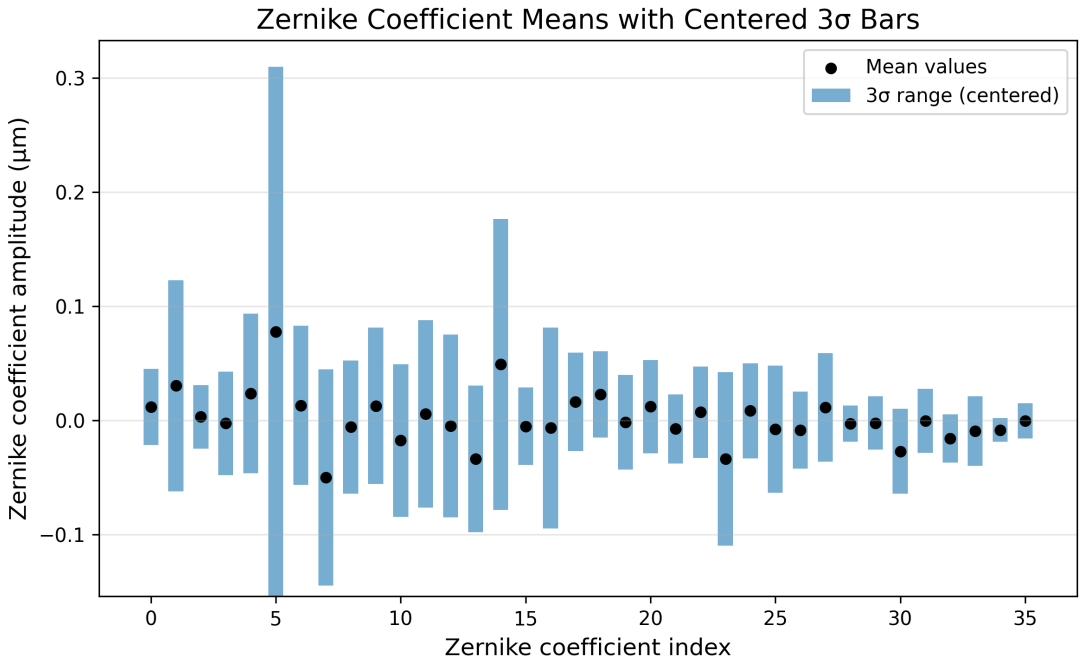

Figure 4.13: Distribution of the 36 Zernike coefficients used to generate new phase maps.

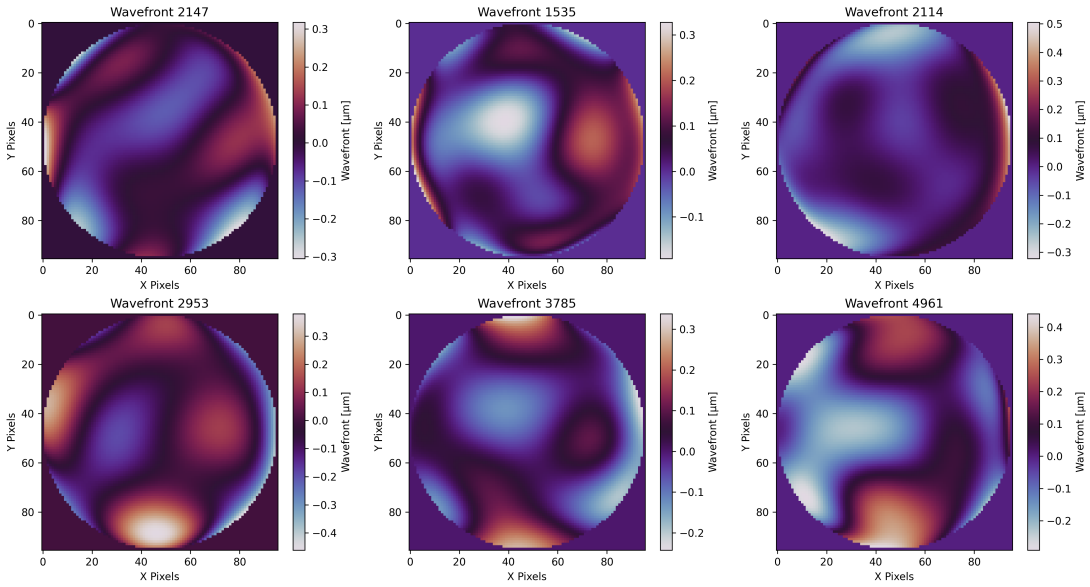

Figure 4.14: Six randomly drawn input phase maps used to train the PINN.

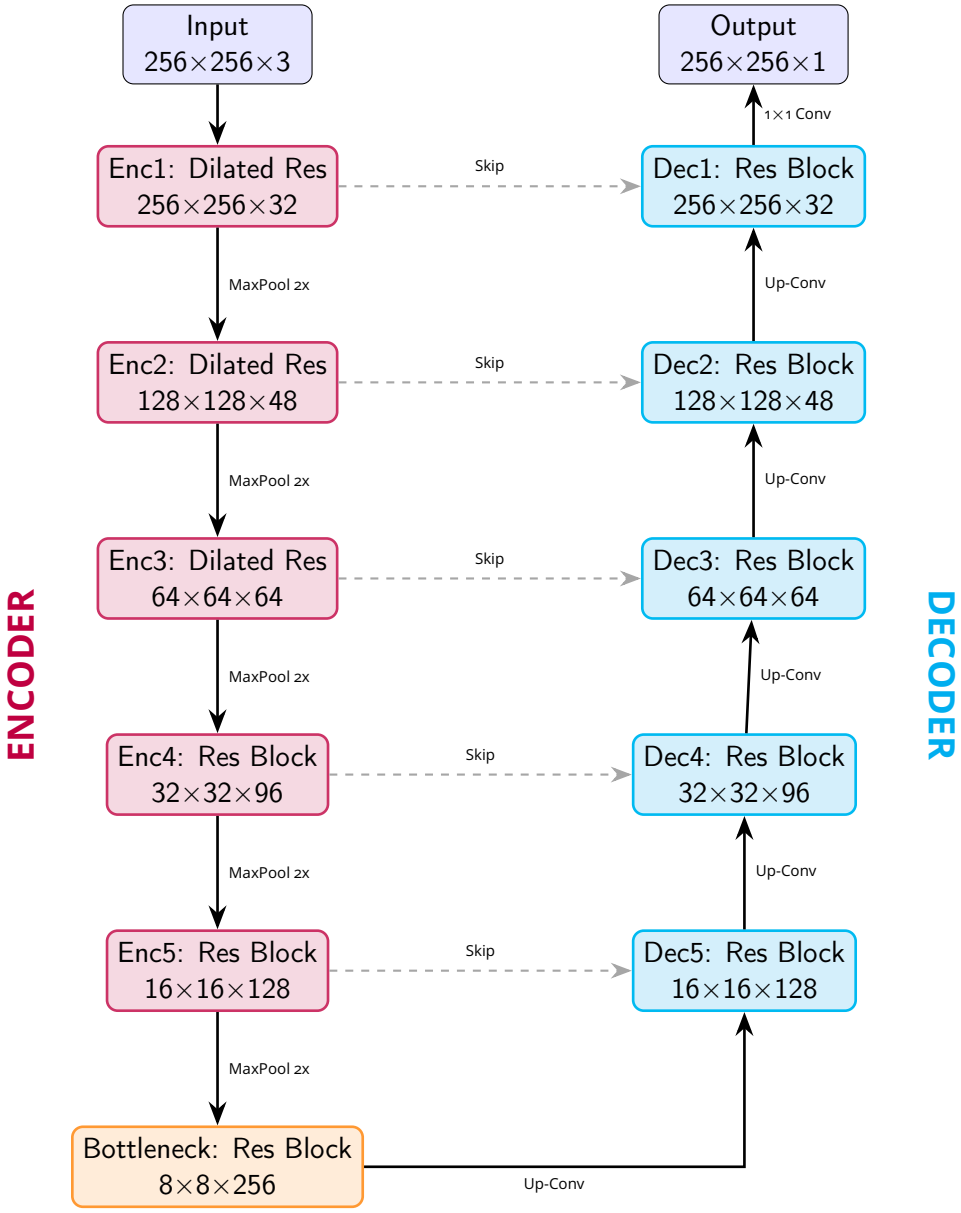

Figure 4.15: Schematic representation of the U-Net architecture used for phase retrieval. In the present case, the three input amplitude planes correspond to $z = -2$ mm, $z = 0$ mm, and $z = +2$ mm. The output is the phase map corresponding to the focus at $z = 0$ mm.

block, which itself contains two 2D convolution layers with a ReLU in between, and a skip connection that adds the block's input to its output before applying a final ReLU. Some residual blocks use dilated convolutions to increase the receptive field, allowing the network to capture diffraction features over multiple spatial scales. After each layer, a max pooling operation downsamples the feature maps, while the number of channels increases to capture progressively more abstract representations. This process is repeated over five layers until the bottleneck is reached.

- **Bottleneck:** No downsampling is performed in the bottleneck. However, the number of feature maps is doubled to increase the network's representational capacity.
- **Decoder:** The decoder mirrors the encoder but uses upsampling operations to progressively restore the spatial resolution. At each decoder layer, the encoder's corresponding feature maps are concatenated via skip connections to preserve fine-grained spatial information.
- **Output:** The network outputs a single phase map corresponding to the focused plane.

Skip Connections: Skip connections are employed to directly transfer high-resolution features from the encoder to the decoder. This helps the network retain fine spatial details that might otherwise be lost during downsampling, improving the reconstruction of high-frequency structures in the phase map.

4.7.3 Loss function

We tested several supervised loss functions inspired by the work of [9]. The three supervised losses explored were:

$$\mathcal{L}_a = 1 - \frac{1}{N} \sum \frac{\Phi_1 \cdot \Phi_2}{\max(\|\Phi_1\|_2, \epsilon)\ \max(\|\Phi_2\|_2, \epsilon)} \tag{1}$$

$$\mathcal{L}_b = \frac{1}{N} \sum (\Delta\Phi)^2 \tag{2}$$

$$\Delta\Phi = (\Phi_2 - \Phi_1 + \pi) \bmod (2\pi) - \pi \tag{3}$$

$$\mathcal{L}_c = 1 - \frac{1}{N} \left\| \sum e^{j\Phi_1} e^{-j\Phi_2} \right\|_2 \tag{4}$$

where Φ_1 denotes the predicted phase and Φ_2 the true focus-plane phase.

In addition to these supervised losses, we investigated an unsupervised physics-based loss, which yielded significantly better performance than any of the three losses above. This loss enforces physical consistency by propagating the predicted focus field to the adjacent measurement planes and comparing only the amplitudes.

$$\begin{aligned} E_{\text{focus}} &= A_{\text{focus}}\, e^{j\Phi_{\text{pred}}} \\ E_{+} &= P_{z_1}\big[E_{\text{focus}}\big], \qquad E_{-} = P_{z_2}\big[E_{\text{focus}}\big] \\ \mathcal{L}_{\text{prop}} &= \frac{1}{2}\left(\frac{\|\,|E_{+}| - |A_{+}|\,\|^2}{\|A_{+}\|^2} + \frac{\|\,|E_{-}| - |A_{-}|\,\|^2}{\|A_{-}\|^2}\right) \end{aligned} \tag{4.22}$$

Here, P_{z_1} and P_{z_2} denote propagation to the forward and backward defocused planes, respectively. The quantities A_{+} and A_{-} are the measured amplitudes at these two planes, while A_{focus} is the amplitude at the nominal focus plane. Propagation is implemented using the angular spectrum decomposition method. In Fig. 4.16 we present the loss $\mathcal{L}_{\text{prop}}$ monitored over 200 epochs.

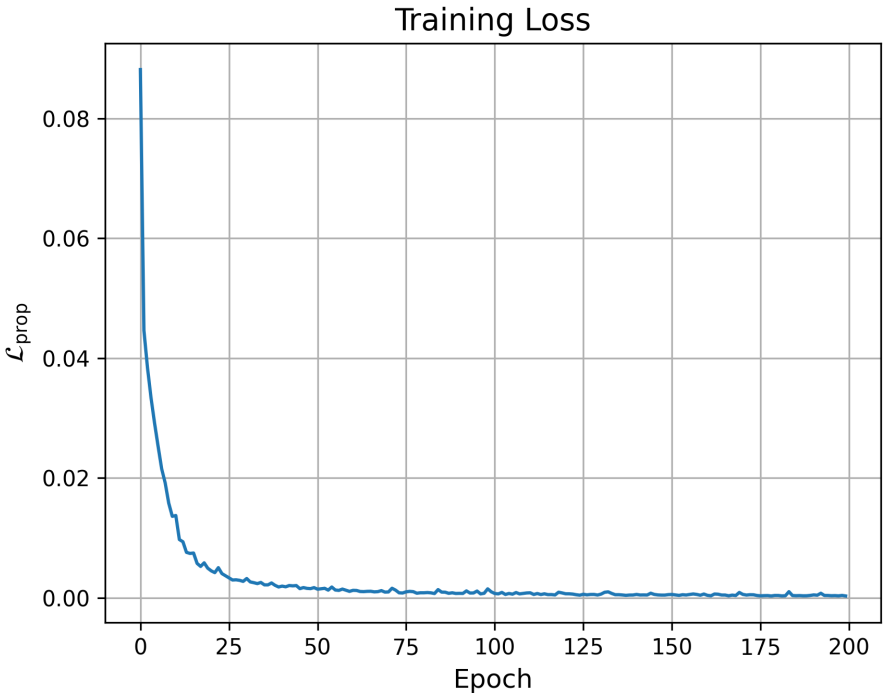

Figure 4.16: Training loss monitored over 200 epochs.

In Fig. 4.17, we show the focal-plane amplitude for a randomly selected test sample. Using this amplitude, we then propagate the field to the forward plane using the phases predicted by the different reconstruction methods. The results of this comparison are presented in Fig. 4.18.

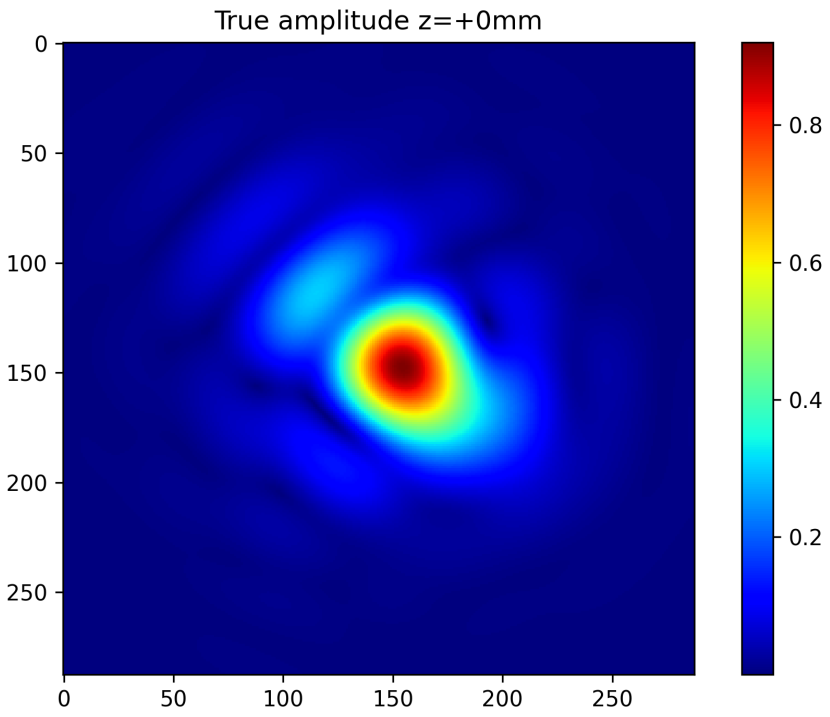

Figure 4.17: Focal-plane amplitude for a randomly chosen test example. The corresponding propagated fields are shown in Fig. 4.18.

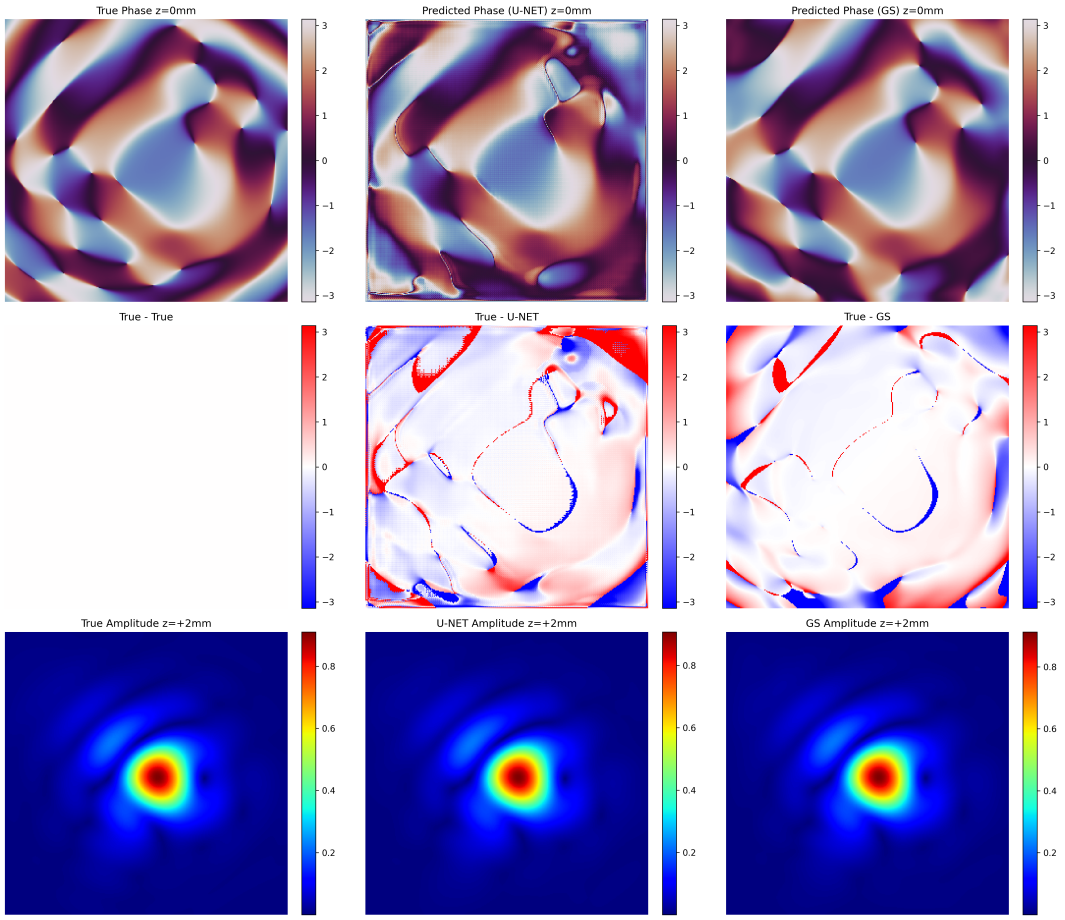

Figure 4.18: Phase retrieval performance of the GSA and U-Net for a randomly chosen test example. **Top row:** True phase at focus (left), U-Net prediction (center), and GS prediction (right). **Middle row:** Phase differences. From left to right: true–true (zero everywhere), true–U-Net, and true–GS. Both methods show negligible error in the central region, where most of the pulse energy is concentrated. **Bottom row:** Propagated amplitude in the forward plane. Both methods reproduce the true beam with high accuracy; the U-Net and GS achieve relative amplitude-squared errors of 1.9×10^{-4} and 4×10^{-5}, respectively.

4.7.4 Conclusion

In conclusion, we developed a physics-informed deep learning framework for phase retrieval that achieves performance comparable to that of the Gerchberg–Saxton algorithm while dramatically reducing computational time. Unlike iterative reconstruction methods, the proposed U-Net architecture enables near–real-time inference and directly predicts the unwrapped phase, eliminating the need for post-processing steps.

The incorporation of a physics-based propagation loss was essential for achieving robust, physically consistent reconstructions. These results demonstrate that the proposed approach is a strong candidate for real-time diagnostic applications and closed-loop feedback phase control of ultrafast laser systems.

This algorithm could eventually be applied to experimental data, where the operator would intentionally generate a large number of controlled aberrations, using, for example, a deformable mirror,to construct a representative training dataset.

References

[1] Christophe Dorrer and Sae-Woo Bahk. “Spatio-spectral characterization of broadband fields using multispectral imaging”. In: *Optics Express* 26.25 (2018), pp. 33387–33399.

[2] Joseph W. Goodman. *Introduction to Fourier Optics*. 2nd ed. New York: McGraw-Hill, 1996.

[3] Kyoji Matsushima and Tomoyoshi Shimobaba. “Band-Limited Angular Spectrum Method for Numerical Simulation of Free-Space Propagation in Far and Near Fields”. In: *Optics Express* 17.22 (2009), pp. 19662–19673. doi: 10.1364/OE.17.019662.

[4] Dominikus Noll. “Alternating projections with applications to Gerchberg–Saxton error reduction”. In: *Set-Valued and Variational Analysis* 29.4 (2021), pp. 771–802. doi: 10.1007/s11228-021-00564-5.

[5] Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational Physics* 378 (2019), pp. 686–707. doi: 10.1016/j.jcp.2018.10.045.

[6] Hamidreza Eivazi, Yuning Wang, and Ricardo Vinuesa. “Physics-informed deep-learning applications to experimental fluid mechanics”. In: *Measurement Science and Technology* 35.7 (2024), p. 075303. doi: 10.1088/1361-6501/ad2f1d.

[7] Kaiqiang Wang et al. “On the use of deep learning for phase recovery”. In: *Light: Science & Applications* 13.1 (2024), p. 4.

[8] Ayan Sinha et al. “Lensless computational imaging through deep learning”. In: *Optica* 4.9 (2017), pp. 1117–1125. doi: 10.1364/OPTICA.4.001117.

[9] Chunjing Chen, Chisheng Wang, and Haidong Yu. “Loss function design for wrapped phase fitting in InSAR deep learning network training”. In: *IEEE Geoscience and Remote Sensing Letters* 21 (2024), pp. 1–5.

5 -Compressor deformation

5.1 Introduction

High-power terawatt (TW) and petawatt (PW) laser systems based on the chirped pulse amplification (CPA) scheme are used in a wide range of scientific fields, including laser-driven particle acceleration, high-energy-density physics, and attosecond science. In these applications, achieving the nominal intensity at the focus is critical. However, deformable mirrors are not able to fully compensate for spatio-temporal coupling (STC) effects, which can have a drastic impact and cannot be detected using standard wavefront sensors or spectral phase detectors such as FROG[1] or SPIDER[2]. STCs are chromatic aberrations of the laser beam, generally defined as the inability to describe the laser electric field as the product of a function of space and a function of time.

Two well-known types of STCs[3] are spatial chirp (SC) and pulse front tilt (PFT). PFT arises from angular dispersion, causing the laser pulse front to tilt relative to the direction of propagation. SC, on the other hand, occurs when different wavelengths within the beam are spatially separated along the transverse direction of the laser.

In laser wakefield acceleration (LWFA), controlling the laser profile is essential not only to manage the ionization injection process and wakefield generation through the intensity at focus, but also to produce high-quality electron beams and ensure stable operation. For instance recent studies have shown the direct link between electron beam steering and PFT[4, 5, 6]

In CPA systems, STCs primarily arise at the optical compressor stage, but they may also result from uncompensated dispersive elements along the laser beam path. For example, misalignment of the optical compressor leads to residual angular dispersion in the near field, which manifests as STCs in the far field [7]. Although fine-tuning of the optical compressor can mitigate these effects in static operation. When the laser system operates at high fluence, heat-induced deformations of the optical compressor gratings can occur, leading to additional STCs [8, 9]. While Leroux et al. characterized the effect of heat-induced deformations on the laser pulse's spatio-spectral properties, they did not account for the additional diffraction introduced by the modified grating geometry. In this section, we extend their approach by first deriving a modified grating equation (MGE) that incorporates this additional diffraction at both gratings. We then present a code based on this equation that calculates the additional optical path length in the compressor due to thermal deformation. Next, we analyze the impact of these deformations on the laser electric field in both the near field and far field using a Zernike polynomial analysis. The results of the code are then compared with experimental measurements acquired with the multi-spectral camera. Finally, in the appendix, we provide further comments on the MGE.

5.2 MGE derivation

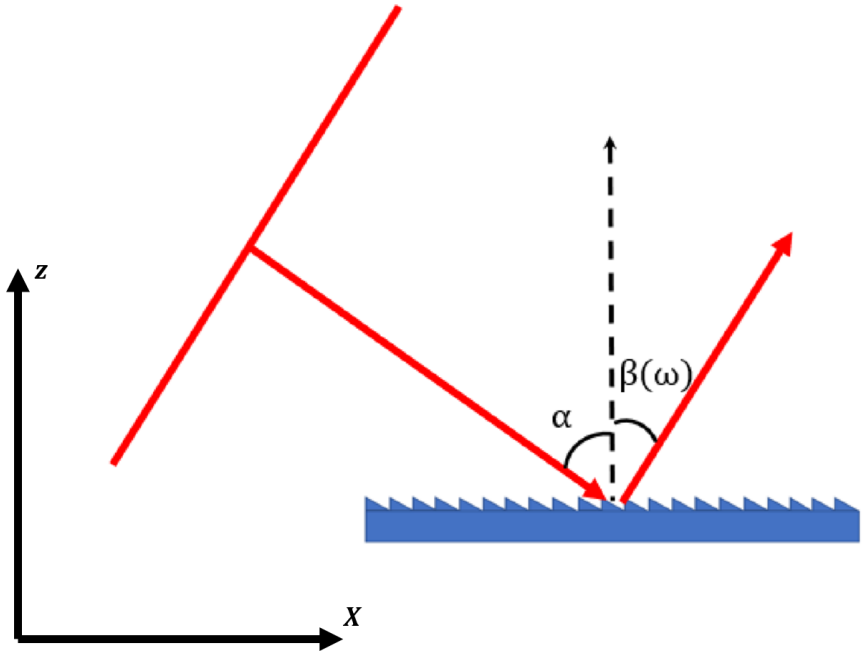

Figure 5.1: Schematic of the deformed grating (blue), incident plane wave, and diffracted rays (red).

To derive the Modified Grating Equation (MGE), we consider the diffraction of a plane wave incident on a deformed reflection grating at an angle of incidence α. We work in two dimensions, within the (x, z) plane. The x-axis is tangential to the undeformed grating surface (along the direction of the grating grooves), while the z-axis is normal to the undeformed surface. The origin of the coordinate system is set at $(x_0, z_0) = (0, 0)$.

When the grating is subjected to thermal stress, it undergoes a local deformation described by a height function $z_1(x)$, representing the displacement of the surface relative to the nominal (undeformed) plane. We now examine the optical path length from the incident plane wave to a point on the deformed grating surface:

$$\mathbf{s}_{\text{in}} \cdot \mathbf{r} = -\sin(\alpha)\, x + \cos(\alpha)\, z_1(x), \tag{5.1}$$

where $\mathbf{s}_{\text{in}}$ is the unit vector along the incident direction, and $\mathbf{r}$ is the position vector of a point on the grating.

To express the far-field amplitude, we use the Fraunhofer diffraction integral. The optical path length from a point on the grating to an observation point $M(X, Z)$ in the far field is

$$r = \sqrt{(X - x)^2 + (Z - z_1(x))^2}. \tag{5.2}$$

For $R = \sqrt{X^2 + Z^2} \gg X, x, z_1(x)$, we can expand r as

$$r \simeq R - \frac{Xx}{R} - \frac{Zz_1(x)}{R} + \frac{x^2 + z_1(x)^2}{2R}. \tag{5.3}$$

Using the Fraunhofer approximation and assuming small deformation amplitudes ($z_1(x)^2 \ll z_1(x)$), we neglect the quadratic terms to obtain

$$r \simeq R - \sin(\beta)\, x - \cos(\beta)\, z_1(x), \tag{5.4}$$

The far-field electric field is then given by

$$E(\beta) \propto \int_{-Nd/2}^{Nd/2} A(x)\, \exp\left[-ik\left((\sin\alpha + \sin\beta)\, x + z_1(x)\,(\cos\beta - \cos\alpha)\right)\right] dx, \tag{5.5}$$

where k is the wavenumber, N is the number of illuminated grooves, d is the grating period, and $A(x)$ is the grating transmission (or reflection) function.

We can express $A(x)$ as a Fourier series:

$$A(x) = \sum_{m=-\infty}^{+\infty} c_m\, e^{i\frac{2\pi m x}{d}}, \tag{5.6}$$

where c_m are the Fourier coefficients corresponding to diffraction order m.

Substituting this into the field integral yields

$$E(\beta) \propto \sum_{m=-\infty}^{+\infty} c_m \int_{-Nd/2}^{Nd/2} \exp\left\{-ik\left[(\sin\alpha + \sin\beta)\, x + z_1(x)\,(\cos\beta - \cos\alpha)\right] + i\frac{2\pi m x}{d}\right\} dx. \tag{5.7}$$

The diffracted angles correspond to directions of constructive interference, determined by the *stationary phase condition* [10] (appendix 3 page 888) :

$$\frac{\partial \Phi}{\partial x} = 0, \tag{5.8}$$

where Φ denotes the total phase of the exponential. Evaluating this condition gives

$$\sin(\alpha) + \sin(\beta) + \frac{dz_1(x)}{dx}\,(\cos\beta - \cos\alpha) = \frac{m\lambda}{d}. \tag{5.9}$$

This is the **Modified Grating Equation (MGE)**, which reduces to the classical grating law when $z_1'(x) = 0$ (flat surface):

$$\sin(\alpha) + \sin(\beta) = \frac{m\lambda}{d}. \tag{5.10}$$

5.3 Code structure

The code presented here calculates the impact of compressor deformation, due to thermal stress, caused by the laser pulse electric field. It is based on a ray-tracing method, where the optical path length resulting from the deformation is accounted for by considering both the physical displacement of the grating surface and the modified diffraction angle.

The first stage is to quantify the physical displacement from the nominal, undeformed grating configuration using COMSOL MULTIPHYSICS. The rays are then diffracted at the first grating ($G1$) using the Modified Grating Equation (MGE), and the total round-trip path length from $G1$ back to the output plane is calculated using a geometric analysis. After this, the field is propagated into the far field, where further analysis is performed. The following subsections detail each stage of the code.

5.3.1 Deformation calculation Using COMSOL MULTIPHYSICS

Heat-induced deformations are simulated using COMSOL MULTIPHYSICS, a physics-based solver that computes field quantities by numerically solving differential equations. The user must specify the geometry, material properties, and the relevant modules, such as the heat module for solving the heat equation. Mesh resolution must also be defined, along with input parameters such as the laser intensity.

In this study, COMSOL first solves the heat equation for the grating pair used in the LASERIX facility. The two gratings are represented as rectangular parallelepipeds matching the real gratings' size. Only the substrate (PYREX, borosilicate) is modeled, as it is expected to dominate the thermal response. Furthermore, grating specifications are unavailable, as they were purchased 20 years ago.

A rectangular mesh is used for both gratings due to the geometry of the problem. The mesh resolution is chosen to be fine, as the computation time in COMSOL is relatively short (a few seconds). The simulation is temporal, spanning from 0 to 70 minutes.

The laser surface power in units of [W/m²] is defined as follows:

$$Q(x,y) = 2P(1-R)\frac{1}{\pi\sigma_x\sigma_y}\exp\left[-2\left(\frac{x^2}{\sigma_x^2}+\frac{(y-y_0)^2}{\sigma_y^2}\right)\right] \tag{5.11}$$

Here, P is the laser power in watts (25 W in this study), R is the reflectivity of the gold coating, σ_x and σ_y are the transverse beam sizes on the grating, and $(0, y_0)$ is the center of the laser beam on the grating. The factor $1/(\pi\sigma_x\sigma_y)$ normalizes the intensity such that the integral over the grating surface equals P.

In this study, only the laser average power at 10 Hz is considered, as the temporal response of the grating material over tens of minutes to an hour is insensitive to the laser peak power, responding primarily to the average power.

The reflectivity R is chosen based on the theoretical reflectivity of gold for an s-polarized laser at an incidence angle of 56° and a central wavelength of 810 nm. The theoretical values range from 0.958 to 0.968 depending on the source; we adopt $R = 0.968$.

The values of σ_x and σ_y depend on the grating: for $G2$ (the diffracted beam), σ_x is larger than for $G1$. The laser power is finally scaled by the diffraction efficiency. For $G1$, we have:

$$q_1 = Q_1(y_{\text{top}}) + Q_1(y_{\text{bottom}}) \cdot 0.9^3,$$

and for $G2$:

$$q_2 = Q_2(y_{\text{top}}) \cdot 0.9 + Q_2(y_{\text{bottom}}) \cdot 0.9^2,$$

where 0.9 corresponds to the diffraction efficiency in the -1 order of the grating law, and Q_1, Q_2 are the surfacic powers for $G1$ and $G2$ respectively, see (FIG. 5.2)

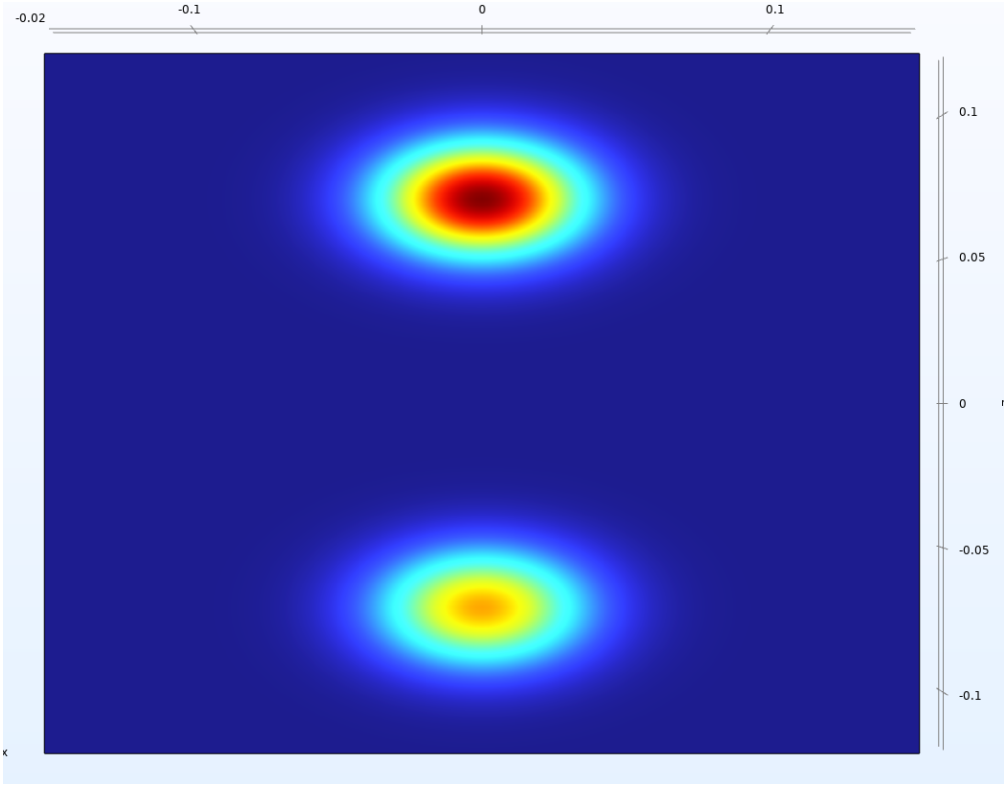

Figure 5.2: Illumination pattern on the first grating $G1$, corresponding to q_1 above.

COMSOL then solves the heat equation using the specified material parameters. The

user only needs to specify the boundary conditions. For $G1$, the boundary conditions are:

$$\begin{cases} -\mathbf{n}\cdot\mathbf{q} = q_1, & S1 \\ -\mathbf{n}\cdot\mathbf{q} = 0, & S2 \\ -\mathbf{n}\cdot\mathbf{q} = 0, & S3 \\ -\mathbf{n}\cdot\mathbf{q} = h_1(T_{\text{ext}} - T), & S4 \end{cases}$$

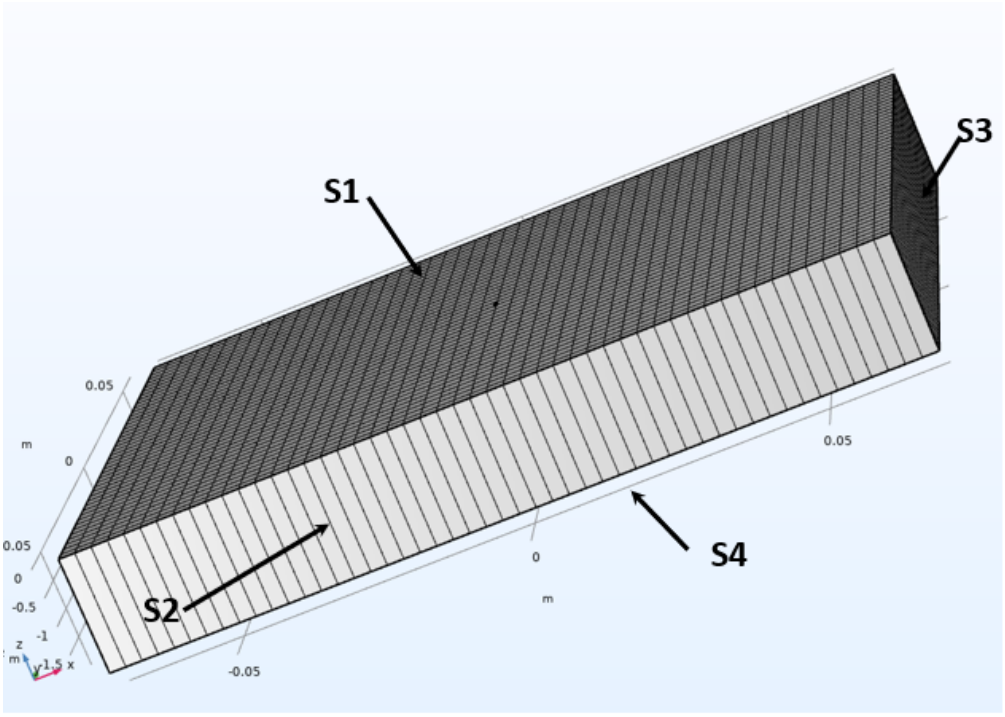

Figure 5.3: Mesh of grating $G1$ with surfaces where the boundary conditions are applied.

For $G2$, the same boundary conditions are applied, replacing q_1 with q_2 and h_1 with h_2. Here, h_1 is the convective heat transfer coefficient, estimated from the Pyrex conductivity and the distance from the front to the back surface, assuming the back surface is in contact with an infinite mount. Surfaces $S2$ and $S3$ use insulating boundary conditions, as they are not in contact with a mount and radiative cooling is negligible.

COMSOL also solves the Cauchy equation of motion to quantify the medium deformation. The initial condition is $u = 0$, and the boundary condition is $u = 0$ on surfaces $S2$ and $S3$. This boundary condition is justified because the displacement at the borders is negligible and provides a fixed region for evaluating vertical displacement which drastically reduce computation time.

With this setup, COMSOL results show a transient response followed by saturation over roughly 5 hour for both temperature and vertical displacement as seen in FIG. 5.4. The vertical displacement is then exported from COMSOL and fed into the propagation code.

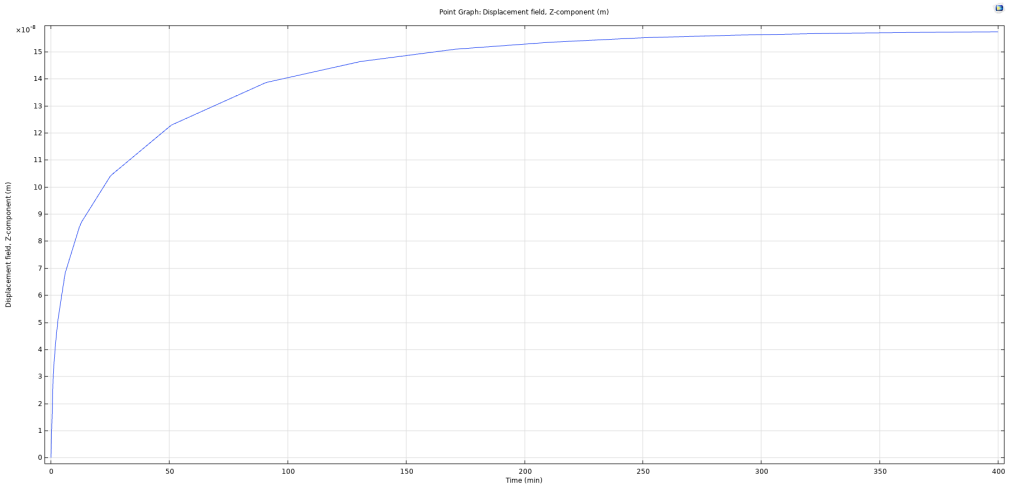

(a) Temporal evolution of the Vertical displacement at the center of the first grating $G1$.

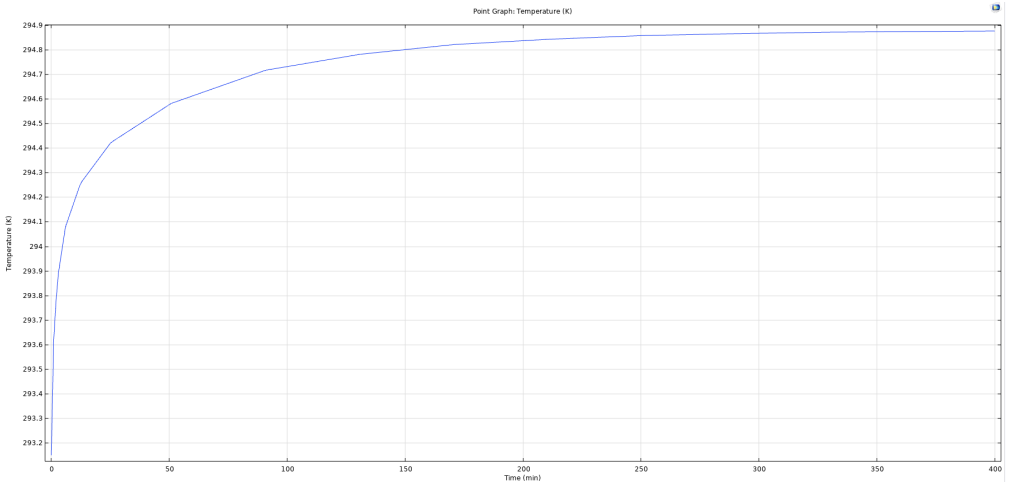

(b) Temporal evolution of the temperature at the center of the first grating $G1$.

Figure 5.4: COMSOL: Temporal evolution of the vertical displacement and temperature for an average laser power of 25 W.

5.3.2 Deformed compressor code

The structure of the propagation code follows the steps illustrated in Fig. 5.5:

- Diffraction at $G1$
- Path length calculation between $G1$ and $G2$
- Diffraction at $G2$
- Reflection at the roof mirror
- Diffraction at $G2$
- Path length calculation between $G2$ and $G1$
- Diffraction at $G1$
- Path length calculation from $G1$ to the output plane

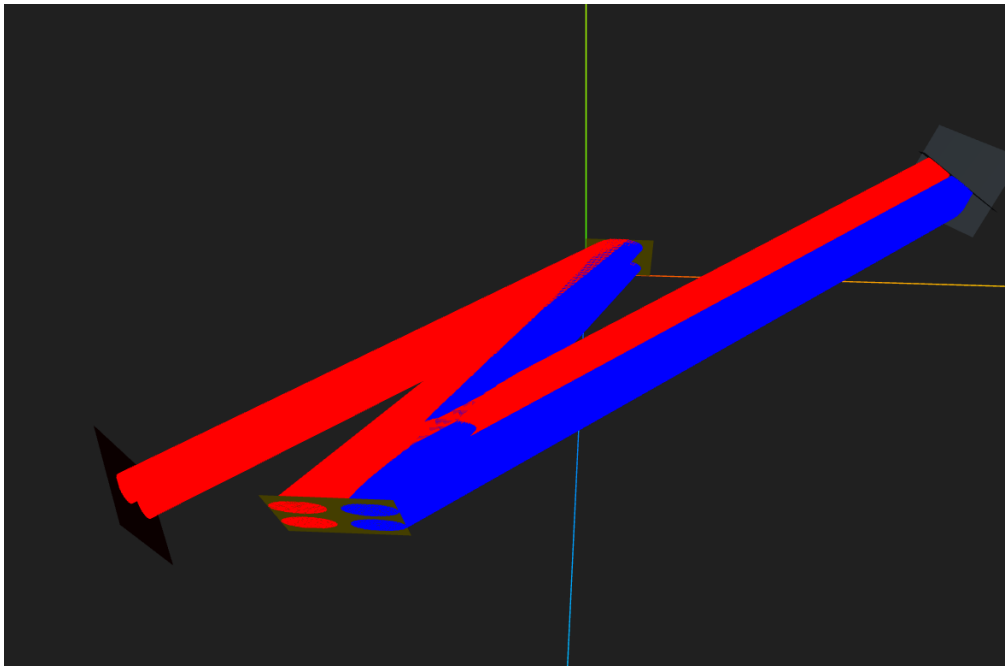

Figure 5.5: 3D rendering of the compressor in the nominal undeformed configuration. Blue rays correspond to a wavelength of 780 nm, and red rays to 836 nm. The output plane is on the left, the two gratings are in the center, and the roof mirror is on the right. The code calculates the optical path length from $G1$ back to the output plane.

5.3.3 Ray propagation at deformed gratings

Step 1: Diffraction at $G1$ using the MGE To perform diffraction at $G1$, we apply the Modified Grating Equation (MGE) while accounting for the three-dimensional geometry of the grating, including the y direction. The local normal and tangential directions on the deformed surface are computed from the gradient of the height function $h = z - z_1(x, y) = 0$ (see Fig. 5.6).

The local normal vector is

$$\mathbf{n} = -\frac{\partial z_1}{\partial x}\,\mathbf{u_x} - \frac{\partial z_1}{\partial y}\,\mathbf{u_y} + \mathbf{u_z},$$

which is normalized to a unit vector:

$$\mathbf{n} \equiv \frac{\mathbf{n}}{\sqrt{\left(\frac{\partial z_1}{\partial x}\right)^2 + \left(\frac{\partial z_1}{\partial y}\right)^2 + 1}}.$$

The local tangential directions are defined as

$$\begin{cases} \mathbf{t_0} = \mathbf{u_x} - (\mathbf{n}\cdot\mathbf{u_x})\,\mathbf{n}, \\ \mathbf{t_1} = \mathbf{u_y} - (\mathbf{n}\cdot\mathbf{u_y})\,\mathbf{n}, \\ \mathbf{t_0} \equiv \dfrac{\mathbf{t_0}}{\sqrt{t_{0x}^2 + t_{0y}^2 + t_{0z}^2}}, \\ \mathbf{t_1} \equiv \dfrac{\mathbf{t_1}}{\sqrt{t_{1x}^2 + t_{1y}^2 + t_{1z}^2}}, \end{cases}$$

where t_{0i} and t_{1i} ($i \in \{x, y, z\}$) are the components along x, y, and z, respectively. Using this local frame, the global output vector is

$$\mathbf{s_{out}} = \mathbf{s_{MGE}} + (\mathbf{s_{in}}\cdot\mathbf{t_1})\,\mathbf{t_1}, \tag{5.12}$$

with the incident vector

$$\mathbf{s_{in}} = (-\sin\alpha, 0, \cos\alpha),$$

and

$$\mathbf{s_{MGE}} = -\sin\beta_{\text{loc}}\,\mathbf{t_0} + \cos\beta_{\text{loc}}\,\mathbf{n}, \tag{5.13}$$

where β_{loc} is obtained by solving the MGE for the incident angle α. Since $\mathbf{s_{out}}$ is defined relative to the local normal $\mathbf{n}$, whereas the MGE gives angles relative to the nominal normal $\mathbf{u_z}$, we convert as

$$\beta_{\text{loc}} = \beta + \arctan\left(\frac{\partial z_1}{\partial x}\right).$$

(See Appendix B for a detailed derivation.)

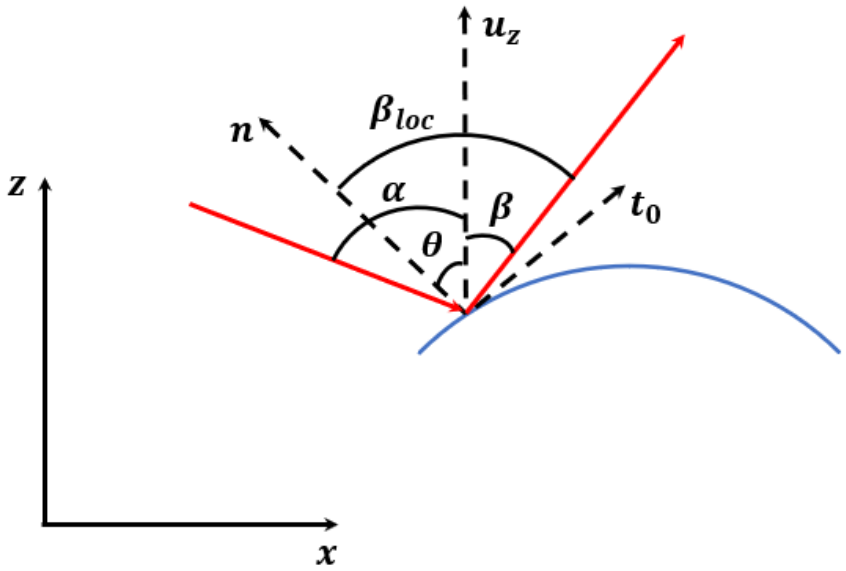

Figure 5.6: Deformed grating surface with local frame vectors $\mathbf{n}$ and $\mathbf{t_0}$. Here, $\theta = \arctan\left(\frac{\partial z_1}{\partial x}\right)$.

Step 2: Path Length Between $G1$ and $G2$ The ray path between $G1$ and $G2$ is parametrized by t:

$$\begin{cases} x_2(t) = x_1 + t\, s_{x_{out}}, \\ y_2(t) = y_1 + t\, s_{y_{out}}, \\ z_2(t) = z_1(x_1, y_1) + t\, s_{z_{out}}, \end{cases}$$

where $z_2(t)$ can also be expressed in terms of $G2$'s surface:

$$z_2(t) = L + z_2(x_2(t), y_2(t)),$$

with L the nominal distance between undeformed gratings, and $z_2(x_2, y_2) < 0$ the deformation of $G2$. This defines the nonlinear function

$$f(t) = z_1(x_1, y_1) + t\, s_{z_{out}} - L - z_2(x_2(t), y_2(t)) = 0, \tag{5.14}$$

which is solved using a Newton–Raphson iteration. Expanding $f(t)$ around $t_0 = L/|s_{z_{out}}|$ gives:

$$f(t) \approx f(t_0) + \left.\frac{df}{dt}\right|_{t=t_0} (t - t_0),$$

leading to

$$t_{n+1} = t_n - f(t_n) \left(\left.\frac{df}{dt}\right|_{t=t_n} \right)^{-1},$$

with derivative

$$\frac{df}{dt} = -\frac{\partial z_2}{\partial x}\, s_{x_{out}} - \frac{\partial z_2}{\partial y}\, s_{y_{out}} + s_{z_{out}}.$$

The iteration stops when $|t_{n+1} - t_n| < 10^{-8}$ m, which typically occurs in about five iterations. This tolerance is appropriate because optical path length variations due to surface deformation are on the order of 10^{-7} m.

Steps 3, 5, and 7: Diffraction Steps These steps are identical to Step 1, except that the sign conventions of the local normals must be handled carefully.

Step 4: Reflection at the Roof Mirror Step 4 computes the path from $G2$ to the roof mirror, along the path between the two mirrors forming the roof, and back to $G2$. A reference position $\mathbf{r_{ref}}$ on $G2$ and a reference vector $\mathbf{s_{ref}}$ toward the roof mirror are defined at the center of the Gaussian laser intensity (810 nm) on $G2$ for the nominal undeformed gratings ($z_1 = 0, z_2 = 0$):

$$\mathbf{r_{ref}} = (x_{2\text{ref}}, y_{2\text{ref}}, L), \quad \mathbf{s_{ref}} = \mathbf{s_{in}} = (-\sin\alpha, 0, \cos\alpha).$$

The apex point and direction of the roof mirrors are defined as:

$$\mathbf{A_{dir}} = \mathbf{s_{ref}} \times \mathbf{u_y}, \quad \mathbf{A_{point}} = \mathbf{r_{ref}} + L_{\text{roof}}\, \mathbf{s_{ref}} + \mathbf{H_y},$$

where $\mathbf{H_y}$ is the half vertical offset of the reference ray after reflection.L_{roof} is the distance to the roof mirror apex point from the reference on $G2$ for$\mathbf{H_y} = (0, 0, 0)$.

The normals of the two roof mirrors are

$$\mathbf{n_1} = -\cos\frac{\pi}{4}\, \mathbf{u_y} - \sin\frac{\pi}{4}\, \mathbf{s_{ref}}, \quad \mathbf{n_2} = \cos\frac{\pi}{4}\, \mathbf{u_y} - \sin\frac{\pi}{4}\, \mathbf{s_{ref}}.$$

The corresponding path lengths are obtained geometrically:

$$l_1 = \frac{(\mathbf{A_{point}} - \mathbf{r_2}) \cdot \mathbf{n_1}}{\mathbf{s_{2out}} \cdot \mathbf{n_1}}, \quad l_2 = \frac{(\mathbf{A_{point}} - \mathbf{r_{M1}}) \cdot \mathbf{n_2}}{\mathbf{s_{M1}} \cdot \mathbf{n_2}},$$

where l_1 is the path from $G2$ to the first mirror, l_2 is the path between mirrors, $\mathbf{r_2} = (x_2, y_2, L + z_2)$ is the position on $G2$, $\mathbf{s_{2out}}$ the outgoing vector from $G2$, and $\mathbf{r_{M1}}$, $\mathbf{s_{M1}}$ the position and the vector after the first mirror:

$$\mathbf{r_{M1}} = \mathbf{r_2} + l_1 \mathbf{s_{2out}}, \quad \mathbf{s_{M1}} = \mathbf{s_{2out}} - 2(\mathbf{s_{2out}} \cdot \mathbf{n_1})\mathbf{n_1}.$$

The final path l_3 from the roof mirror to $G2$ is computed using the Newton–Raphson scheme from Step 2, with initial guess:

$$t_0 = \frac{(L + z_2) - (\mathbf{r_{M2}} \cdot \mathbf{u_z})}{\mathbf{s_{M2}} \cdot \mathbf{u_z}},$$

where $\mathbf{r_{M2}}$ and $\mathbf{s_{M2}}$ are the position and vector after the second mirror, computed analogously to $\mathbf{r_{M1}}$ and $\mathbf{s_{M1}}$.

Step 8: Path Length to the Output Plane Step 8 is similar to Step 4, but now the reference position is on $G1$ and the output plane, with the nominal distance separating the output plane and $G1$.

Implementation Notes All partial derivatives (see FIG. 5.7) of the height function are numerically computed from the exported COMSOL deformation map. Within the code, the height function and its derivatives are evaluated using a bilinear interpolation scheme.

The code is precompiled to a binary using the `numba` Python library, reducing computation time by a factor of 50. It could be further ported to a GPU for even faster inference.

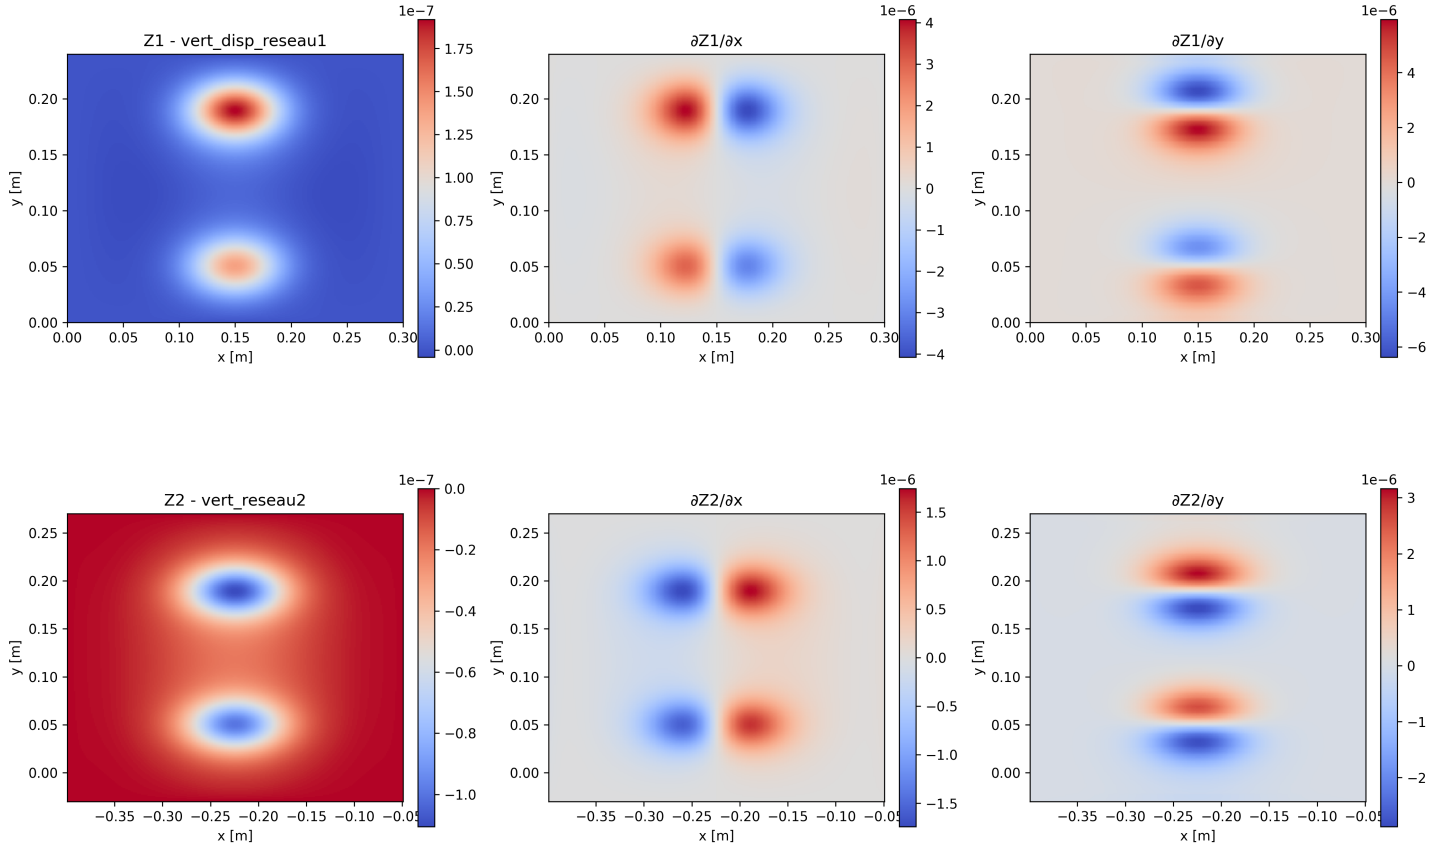

Figure 5.7: Deformation map and partial derivatives for an illumination of 25 W during 1 minute.

5.3.4 Near-field definition and far-field propagation

Once the total optical path length through the compressor is calculated, the phase corresponding to the residual deformation in the output plane is defined as:

$$\Phi(x, y, \omega) = \frac{\omega}{c} \left(OPL_0 - OPL_{\mathrm{STC}} \right), \tag{5.15}$$

where OPL_0 is the optical path length in the nominal, undeformed case, and OPL_{STC} is the corresponding path length in the deformed case.

The complex electric field in the output plane after the compressor is then approximated as:

$$E_{\mathrm{NF,out}} = E_{\mathrm{NF,in}}\, e^{i\Phi(x,y,\omega)}. \tag{5.16}$$

where $E_{\mathrm{NF,out}}$ is defined in the output plane of the compressor illustrated in Fig. 5.5. Equation (5.16) is an approximation: strictly speaking, the output field should be obtained from an integral formulation accounting for the full propagation through the compressor. However, since the output rays intersect the output plane at nearly the same transverse positions as the input near field, a direct phase correction as in Eq. (5.16) is sufficient. In this approximation, the output coordinates relate to the input coordinates as:

$$\begin{cases} x_{\mathrm{out}} \simeq x_{\mathrm{in}}, \\ y_{\mathrm{out}} \simeq y_{\mathrm{in}} - 2\,\mathbf{H_y} \cdot \mathbf{u_y}. \end{cases}$$

Because the output plane makes an angle α with $G1$, we interpolate $E_{\mathrm{NF,out}}$ onto the output grid:

$$\begin{cases} x = x_{gr} \cos\alpha, \\ y = y_{gr}, \end{cases}$$

where x_{gr} and y_{gr} are coordinates on the grating surface.

Finally, the far-field distribution and temporal profile are obtained using a Fourier transform scheme:

$$\begin{aligned} E_{\mathrm{FF}}(k_x, k_y, \omega) &= \mathcal{F}_{x,y}\left[E_{\mathrm{NF,out}}(x, y, \omega)\right], \\ E_{\mathrm{FF}}(k_x, k_y, t) &= \mathcal{F}_{\omega}^{-1}\left[E_{\mathrm{FF}}(k_x, k_y, \omega)\right]. \end{aligned} \tag{5.17}$$

FFT Implementation Notes To ensure accurate far-field calculations and avoid artifacts such as checkerboard phase patterns:

- The near-field grid spacing Δx, Δy must satisfy the Nyquist criterion for the maximum transverse spatial frequency.
- Proper zero-padding should be applied to increase the sampling in Fourier space.
- Use 'fftshift'/'ifftshift' to correctly align the zero-frequency component at the center of the spectrum.

These precautions ensure that the Fourier transform produces a smooth, physically accurate far-field intensity and phase distribution.

5.4 Results

All results presented in this section are obtained either from simulations or from experimental measurements using the multispectral camera (MSCCD) and a spectrometer. We first focus on the measured data. Two distinct measurement sets were performed.

The first measurement corresponds to laser operation at 1 Hz for several minutes, allowing the application of a Gerchberg–Saxton algorithm (GSA) and providing a reference dataset for comparison with measurements performed at higher repetition rate. The second measurement corresponds to laser operation at 10 Hz over a duration of 30 minutes, during which one image was recorded every second. At the end of this acquisition, when the compressor deformation was close to saturation, a GSA was performed using the MSCCD data.

5.4.1 1 Hz laser operation

The measurements were performed on the PALLAS diagnostic table using two instruments: the multispectral camera and a spectrometer. During this measurement, the laser energy was $2.5\ \mathrm{J}$ at the output of the third amplifier, with a repetition rate of 1 Hz. After the compressor gratings, the remaining energy was approximately $1.25\ \mathrm{J}$.

Before performing any measurement, the compressor gratings were carefully aligned using the multispectral camera in order to minimize residual angular dispersion. The deformable mirror was then used to optimize the focal spot. Once these two preparatory steps were completed, a scan along the z direction was performed using the MSCCD, which was mounted on a motorized translation stage.

Images were acquired at three longitudinal positions:

$$z^- = -1.5\ \mathrm{mm}, \quad z = 0\ \mathrm{mm}, \quad z^+ = +1.5\ \mathrm{mm},$$

with $z = 0\ \mathrm{mm}$ defining the focal plane.The Rayleigh length of the beam was roughly $1.6\ \mathrm{mm}$, so these planes sample the near-focus region around the waist. The scan was carried out over several minutes, with ten images recorded at each plane.

This z-scan enabled the application of the Gerchberg–Saxton algorithm to retrieve the spatio-spectral intensity and phase from the three measurement planes. One image out of the ten acquired at each plane was selected as input for the GSA. The results of the measurement and the reconstruction are shown in Fig. 5.8.

As shown in Fig. 5.8, the reconstruction is not perfect. The GSA yields an average relative squared amplitude error, averaged over the six spectral channels, of 8.4% for the backward plane and 3.7% for the forward plane. This discrepancy can be explained by several factors.

First, our problem is ill-posed. Indeed, since we used a single camera to acquire all images, we implicitly assumed that the laser intensity profile remains constant over

several minutes, which is not true in practice; small variations occur from one image to the next.

The second factor is that the laser experienced significant beam-pointing jitter that was not compensated during acquisition. While this does not necessarily prevent convergence of the algorithm, if left uncorrected, the retrieved phase may include spurious contributions due to the jitter. To mitigate this effect, we recentered the images prior to reconstruction. In an ideal setup, one would acquire the images simultaneously on multiple cameras and use a active beam poiting stabilisation with fast steering mirror not yet implemented at the time of this measurement (see chapter 2, LIF module description).

Despite these imperfections, the reconstructed images closely resemble the experimentally measured intensity distributions for both the forward and backward planes across all six spectral channels.

The first important observation from Fig. 5.8 is that, despite optimization with the deformable mirror, the laser clearly exhibits residual spatial aberrations. However, in this configuration, where the compressor is not thermally loaded, the spatio-spectral aberrations remain minimal. Indeed, the intensity patterns across the six spectral channels are very similar, and the measured spatial chirp (SC) is almost negligible.

Since a two-dimensional imaging spectrometer was not available during this measurement, yet an estimate of the spatio-spectral intensity was still required, we developed a dedicated reconstruction algorithm. This algorithm will be particularly useful in the next chapter devoted to realistic particle-in-cell (PIC) simulations, as well as for analysing the 10 Hz measurements.

From the measurement presented in Fig. 5.8, we reconstruct the complex electric field as

$$E(x, y, \omega_i) = \sqrt{I(x, y, \omega_i)}\, e^{i\phi(x,y,\omega_i)}, \tag{5.18}$$

for $i \in [1, 6]$, corresponding to the six measured spectral channels.

Because recentering the images is required to improve GSA convergence and suppress spurious phase contributions, all spatial chirp information is initially removed from the data. To preserve this information, the spatial chirp is reintroduced after the phase retrieval for each spectral channel. This is achieved by first calculating the spatial chirp for each acquired image prior to the GSA, and then reapplying it consistently to both the retrieved phase and intensity. In the 1 Hz case, the spatial chirp is negligible, but this step is crucial for the 10 Hz case where spatial chirp becomes significant. The final electric field is thus written as

$$E(x, y, \omega_i) = \sqrt{I\big(x - SC(\omega_i),\, y - dy(\omega_i),\, \omega_i\big)}\, e^{i\phi\left(x - SC(\omega_i),\, y - dy(\omega_i),\, \omega_i\right)}. \tag{5.19}$$

Here, $SC(\omega_i)$ denotes the spatial chirp along the transverse direction, and $dy(\omega_i)$ represents the centroid displacement along y, with the spectral channel at $\lambda = 803\ \mathrm{nm}$ taken as the reference. The displacement is explicitly denoted dy because, as will be

shown in the 10 Hz case, the spatial chirp predominantly occurs along the x direction, while the displacement along y remains negligible.

After this step, the main limitation is the small number of available spectral channels. The dominant approximation therefore consists in interpolating the electric field from a grid of only six spectral points to a much finer grid of 400 points, such that $i \in [1, 400]$. While this interpolation is not strictly necessary for the present 1 Hz analysis, it becomes essential for the realistic simulations presented in the next chapter.

Finally, the measured spectrum shown in Fig. 5.9 is used to scale the reconstructed intensity. In Fig. 5.9, we also present the reconstructed spatio-spectral intensity distribution, from which it is clear that no measurable spatial chirp is present in the 1 Hz case.

Finally, we present the spatial intensity at focus in the 1 Hz case in Fig. 5.10.

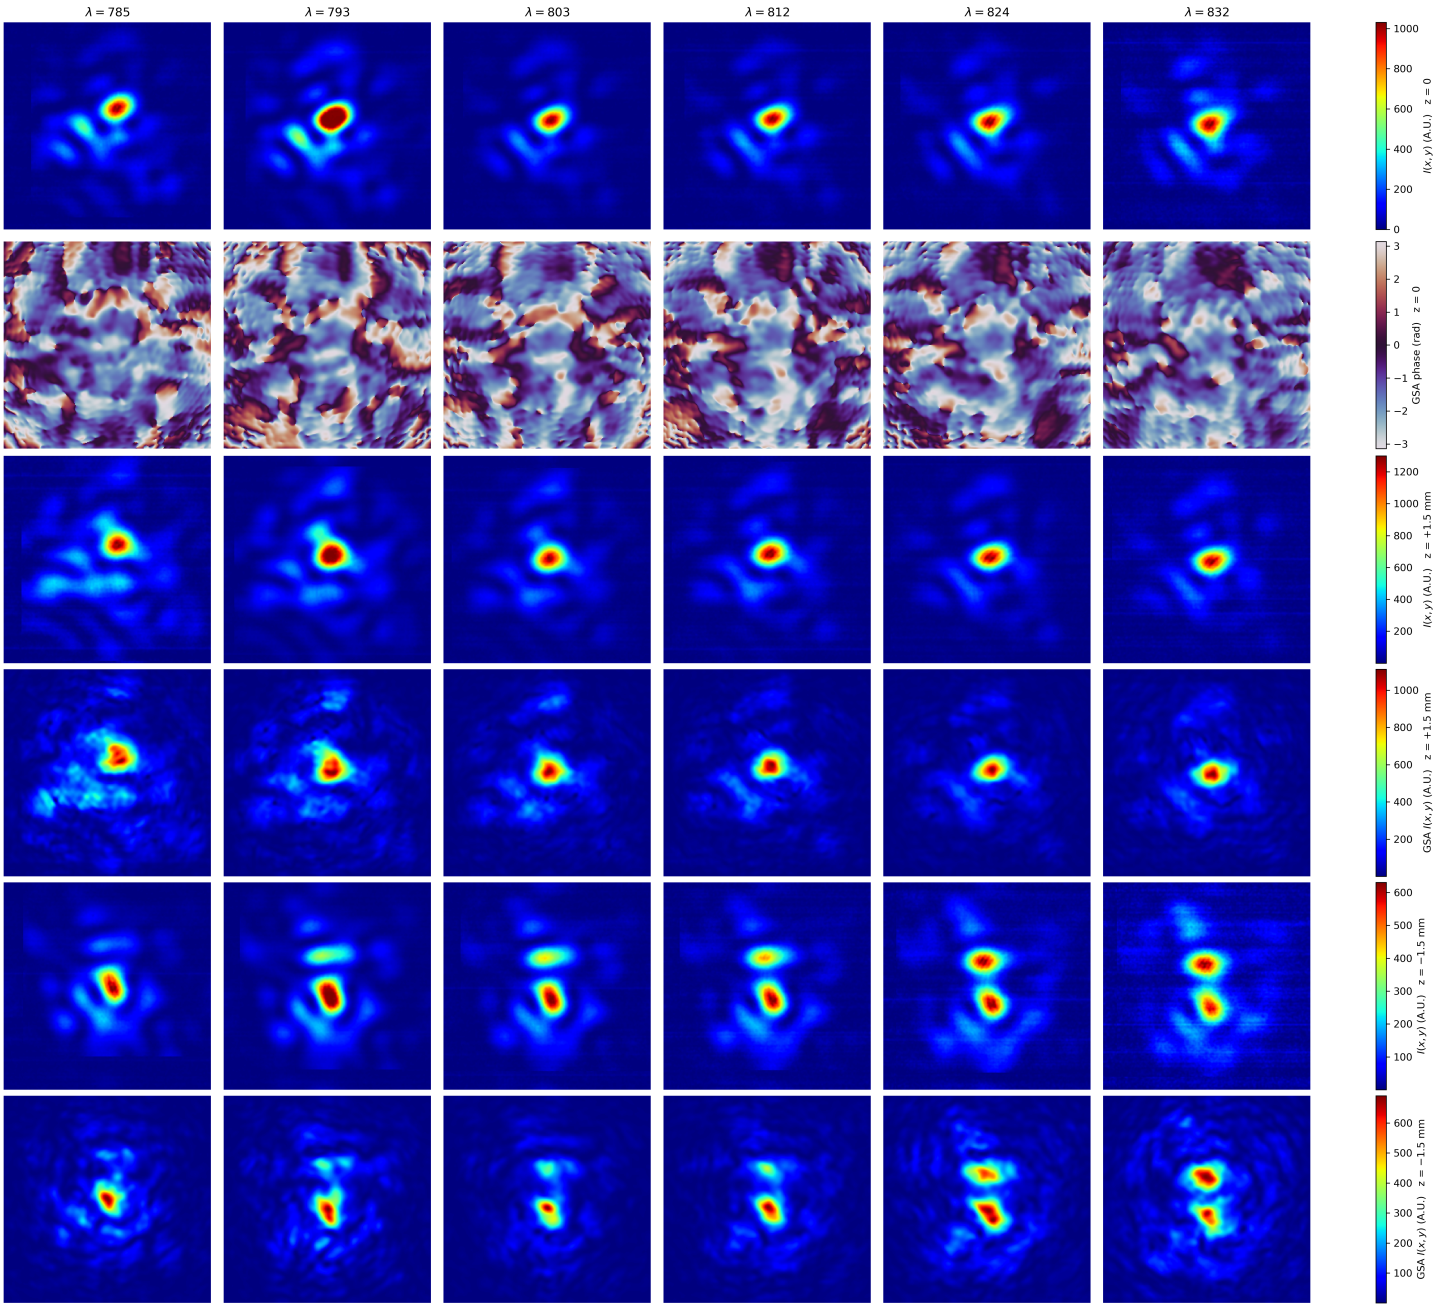

Figure 5.8: Comparison between measured intensity (arbitrary units) and GSA-reconstructed fields for the multispectral camera in the 1 Hz case. Top row: measured intensity at $z = 0\,\mathrm{mm}$. Second row: corresponding spatial phase retrieved by GSA. Third row: measured intensity at $z = +1.5\,\mathrm{mm}$. Fourth row: intensity at $z = +1.5\,\mathrm{mm}$ after propagation using the GSA-retrieved spatial phase. Fifth row: measured intensity at $z = -1.5\,\mathrm{mm}$. Sixth row: intensity at $z = -1.5\,\mathrm{mm}$ after propagation using the GSA-retrieved spatial phase.

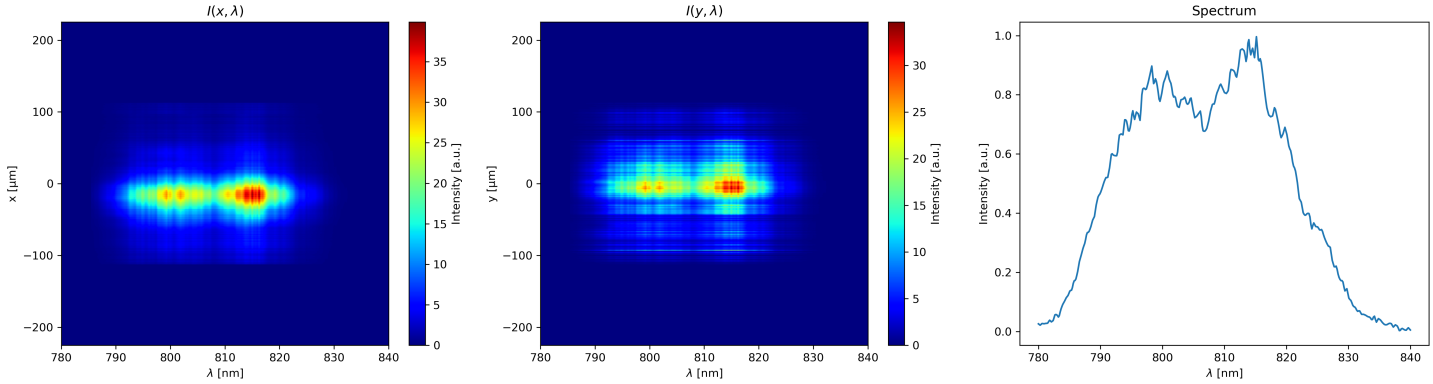

Figure 5.9: Measured spatio-spectral intensity and spectrum for the 1 Hz case.

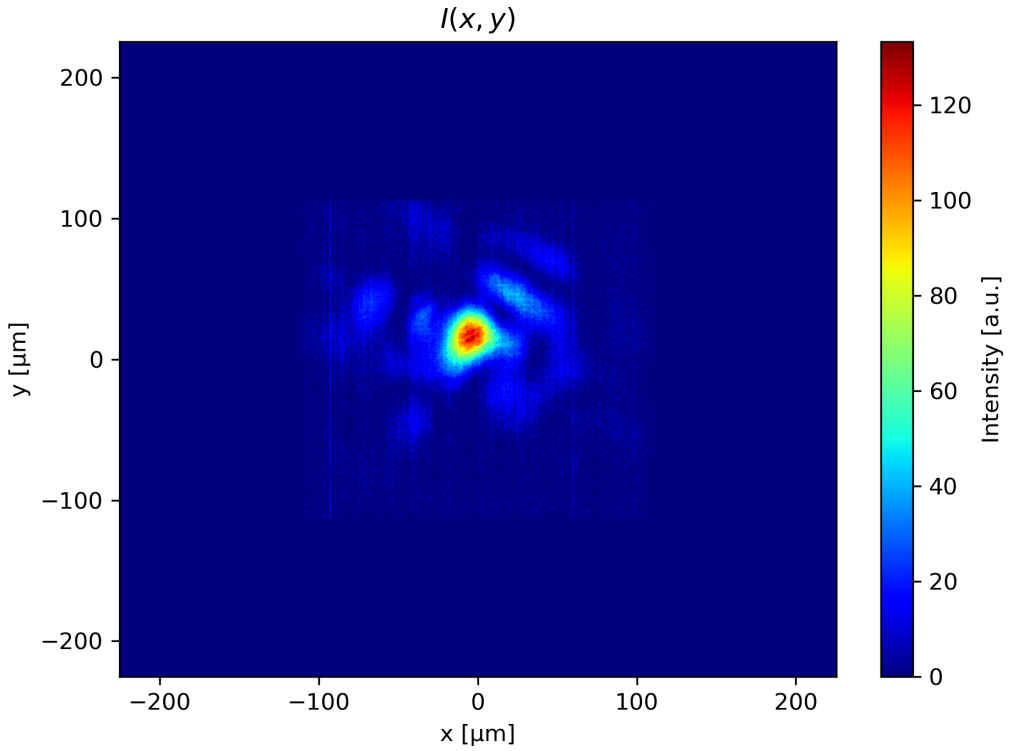

Figure 5.10: Measured spatial intensity in the 1 Hz case. As expected the deformation map resemble the illumination pattern.

5.4.2 10 Hz laser

We now focus on the 10 Hz case. As mentioned previously, images were recorded with the MSCCD over a 30 minute period. At the end of this acquisition, when the spatial chirp ceased to increase, we performed the GSA reconstruction.

In Fig. 5.11, we first present the evolution of the spatial chirp over the 30 minutes of heating, defined as

$$\mathrm{SC} = \sum_{\substack{\lambda \in \text{wavelengths} \\ \lambda \neq \lambda_{\mathrm{ref}}}} |x(\lambda) - x(\lambda_{\mathrm{ref}})| \,, \tag{5.20}$$

where $x(\lambda)$ is the centroid of the image at wavelength λ, and $\lambda_{\mathrm{ref}} = 803\ \mathrm{nm}$ is the reference wavelength.

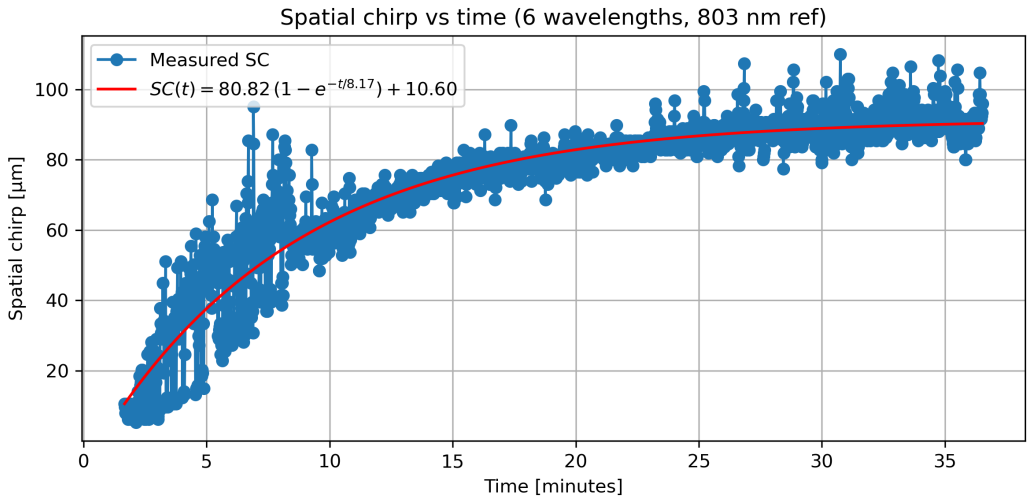

Figure 5.11: Measured spatial chirp as a function of time, defined in Eq. (5.20).

As expected, the temporal evolution of the spatial chirp closely resembles the evolution of the compressor temperature and vertical deformation, showing a curve similar to the charging of a capacitor. Around 30 minutes, the curve reaches saturation, indicating that both the compressor temperature and the vertical deformation have also saturated.

In Fig. 5.12, we present the result of the GSA reconstruction performed after 30 minutes of heating, using the same method as in the 1 Hz case.

The average relative squared amplitude error across channels is 8.7% for the backward plane and 4.9% for the forward plane. Unlike the 1 Hz case, it is immediately apparent that the laser exhibits spatio-spectral coupling (SSC) despite the recentering procedure. Indeed, the intensity pattern varies across channels. Using the same procedure as in the 1 Hz case, we present the reconstructed spatio-spectral intensity and the spatial intensity at focus in Fig. 5.13.

Contrary to the 1 Hz case, some spatial chirp is clearly visible. It occurs predominantly along the x direction, with minor contributions along y.

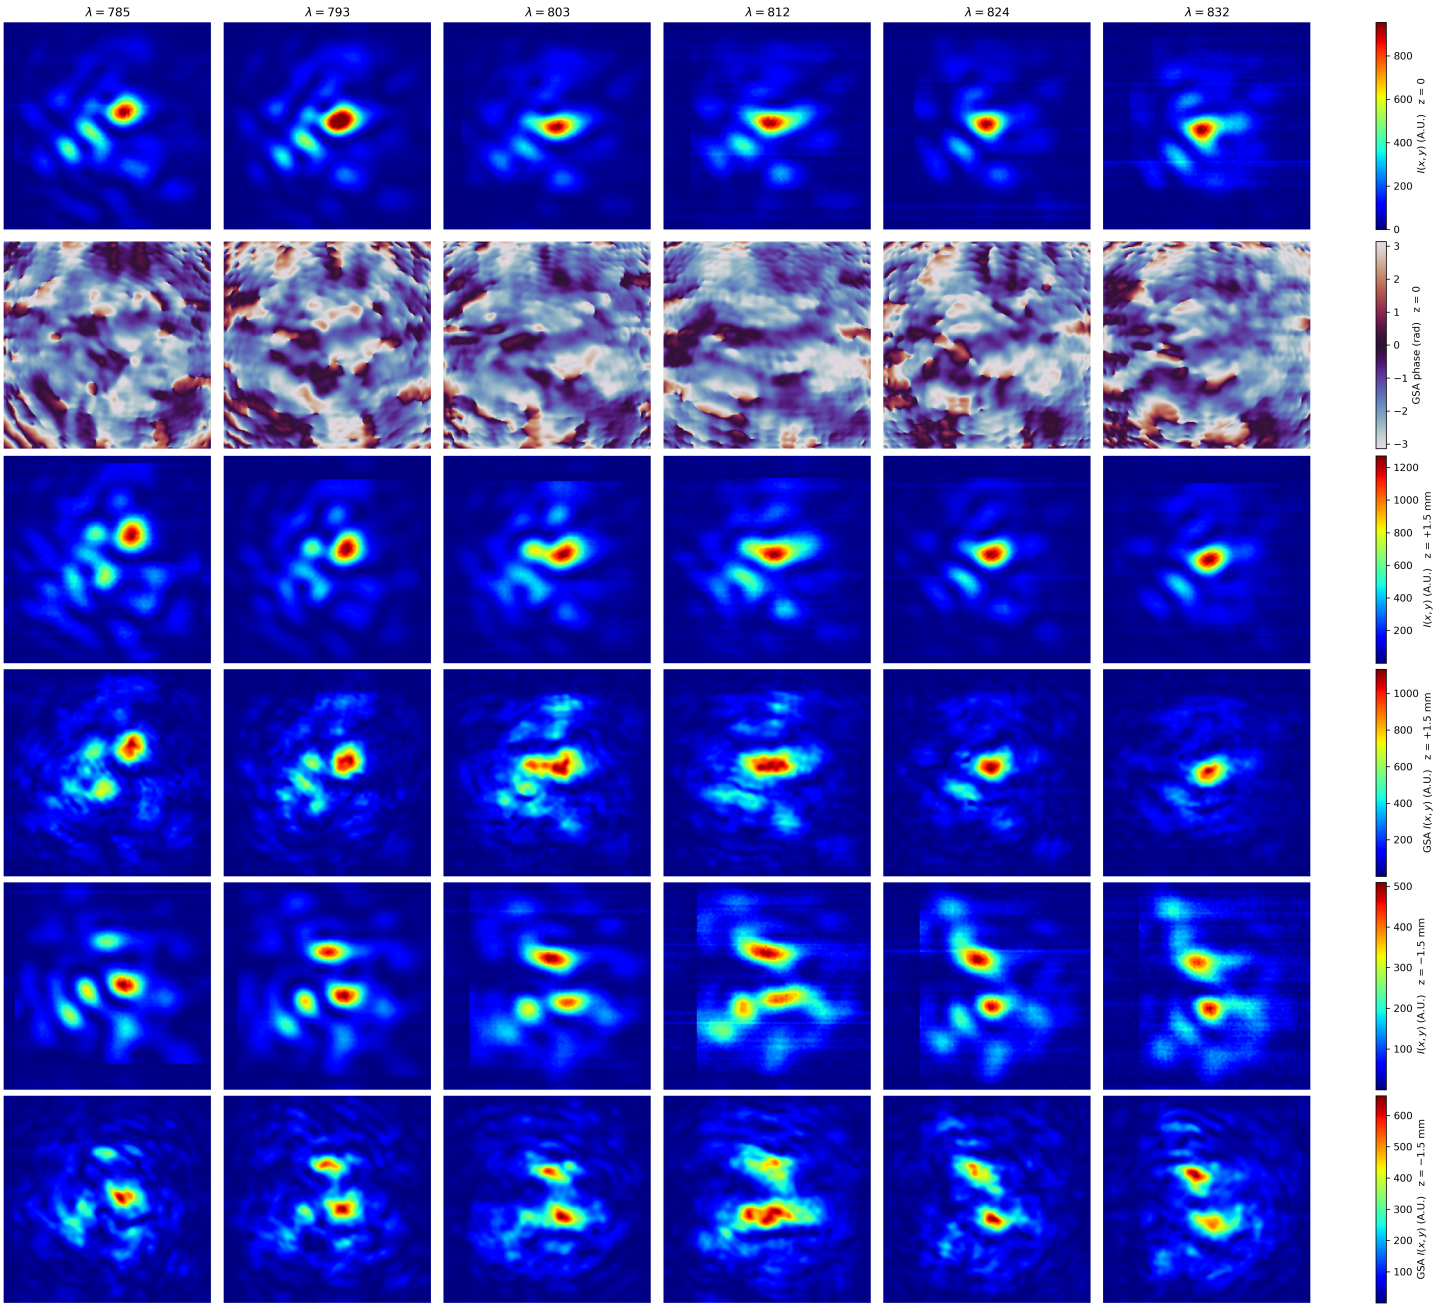

Figure 5.12: Comparison between true intensity (A.U.) and GSA-reconstructed fields for the multi-spectral camera after 30 minutes of heating in the 10 Hz case. Top row: true intensity at $z = 0\,\mathrm{mm}$. Second row: corresponding spatial phase retrieved by GSA. Third row: true intensity at $z = +1.5\,\mathrm{mm}$. Fourth row: intensity at $z = +1.5\,\mathrm{mm}$ after propagation using the GSA spatial phase. Fifth row: true intensity at $z = -1.5\,\mathrm{mm}$. Sixth row: intensity at $z = -1.5\,\mathrm{mm}$ after propagation using the GSA spatial phase.

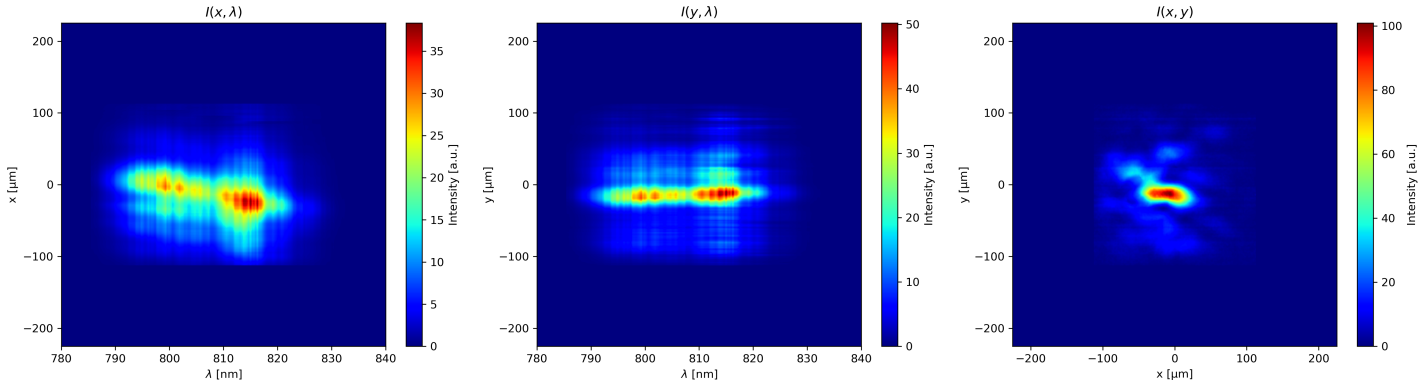

Figure 5.13: Measured spatio-spectral intensity and spatial intensity after 30 minutes of heating in the 10 Hz case.

5.4.3 Simulation results and comparison

We now compare the simulation results with the 10 Hz measurement after 30 minutes of heating. We first note that the simulation results are in better agreement with the experimental data if we consider the simulation output after only 1 minute of heating, using the same parameters as described in the previous sections. This can be explained by several factors.

First, the deformation pattern shown in Fig. 5.7 remains largely unchanged throughout the heating process. In other words, the shape of the deformation and its derivatives does not evolve significantly; only the amplitude scales with heating time.

Second, our model considers only the gold coating reflectivity. In reality, a dielectric layer beneath the gold increases the effective reflectivity seen by the laser, resulting in smaller deformation amplitudes after 30 minutes of heating.

Third, the model assumes a continuous heat source, whereas the laser actually deposits energy in discrete pulses (10 Hz, 10 high-intensity pulses per second), which can slightly alter the heat distribution.

Other potential sources of discrepancy include the estimated convective heat transfer coefficient, the neglect of radiative cooling, and uncertainties in the thermal expansion coefficient and conductivity due to the unknown exact composition of the gratings, which were purchased 20 years ago.

Figure 5.14 presents the simulated output amplitude after the compressor alongside the measured spectrum, corresponding to the measurement shown in Fig. 5.9.

Next, we present in Fig. 5.15 the near-field spatial phase in the same plane as in Fig. 5.5 for the six channels associated with the MSCCD. It is evident that the phase varies across wavelengths even in the near field, due to the compressor deformation.

Figure 5.16 shows the far-field spatial intensity and wrapped phase for six different wavelengths. Both the intensity and phase clearly vary across wavelength, in agreement with the measured results in Fig. 5.12. A pronounced spatial chirp is also observed, predominantly along the x-direction, perpendicular to the grating grooves.

In Fig. 5.17, we present the spatio-spectral intensity along both directions, together with the corresponding spatial intensity. The simulated intensity closely resembles the measured intensity shown in Fig. 5.13. Using equation (5.20), the spatial chirp in the simulation reaches $SC = 134\ \mu\mathrm{m}$, slightly higher than the measured value of $SC = 90\ \mu\mathrm{m}$. This discrepancy can be corrected by scaling the simulated deformation amplitude.

Figure 5.18 shows the spatio-temporal intensity in both directions, revealing a strong pulse-front tilt (PFT) along the x-direction.

Finally, Fig. 5.19 presents the temporal intensity profile. Compared to the nominal undeformed case with $\tau_{\mathrm{FWHM}} = 40\ \mathrm{fs}$, the temporal duration clearly increases due to the compressor deformation. In addition the deformation causes the phase to acquire some group deplay and third order dispersion.

Next, we perform a Zernike decomposition analysis to identify which aberrations

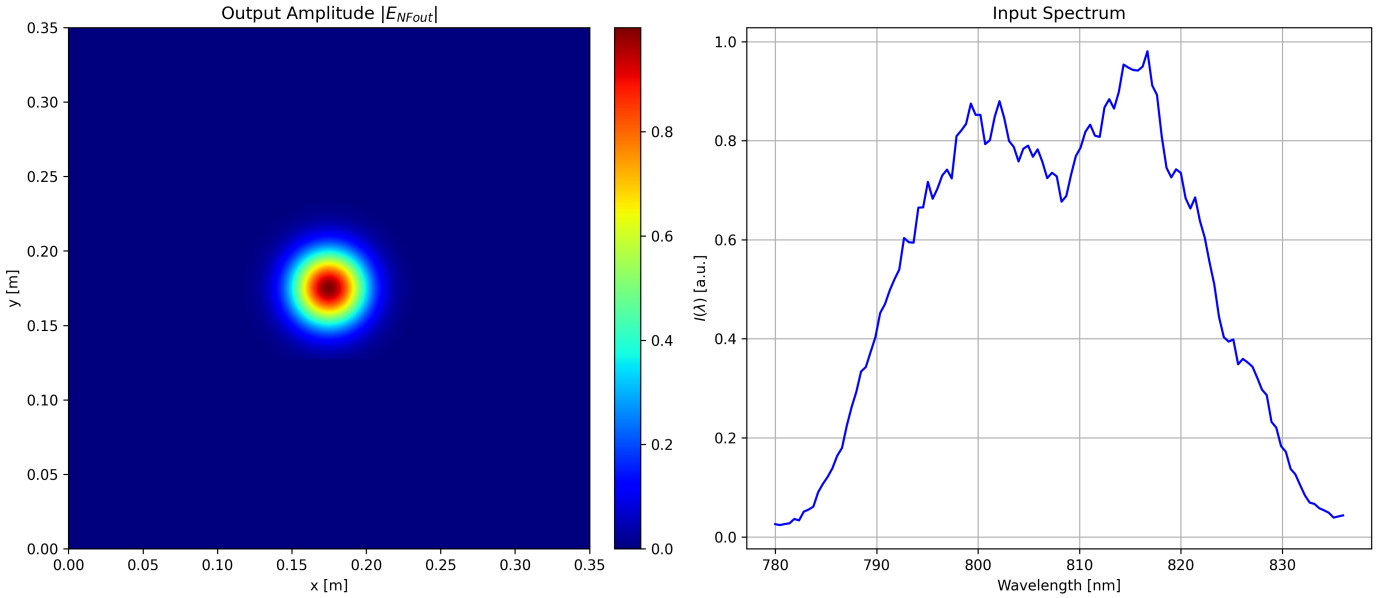

Figure 5.14: Simulation: Near-field amplitude and spectrum in the output plane of Fig. 5.5.

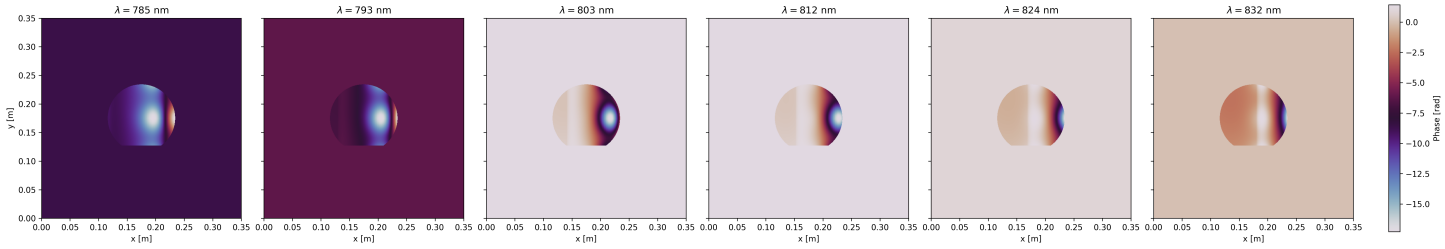

Figure 5.15: Simulation: Near-field spatial phase in the output plane of Fig. 5.5 for different wavelengths after 1 minute of heating.

impact which wavelengths at focus. In Fig. 5.20, we present the amplitude and phase (both wrapped and unwrapped) at the wavelength $\lambda = 803\ \mathrm{nm}$. Fig. 5.21 shows the corresponding Zernike decomposition results for this wavelength. The reconstruction is fairly accurate, with a coefficient of determination $R^2 = 0.92$.

Using the same procedure, we perform the decomposition for every wavelength in the simulation and present the first 12 Zernike polynomials for each wavelength in Fig. 5.22. Each blue point represents a different wavelength. The figure clearly shows that different wavelengths experience different aberrations.

For example, considering the tilt-x aberration, there is a clear slope with a jump around the central wavelength $\lambda = 810\ \mathrm{nm}$. This indicates the presence of wavefront

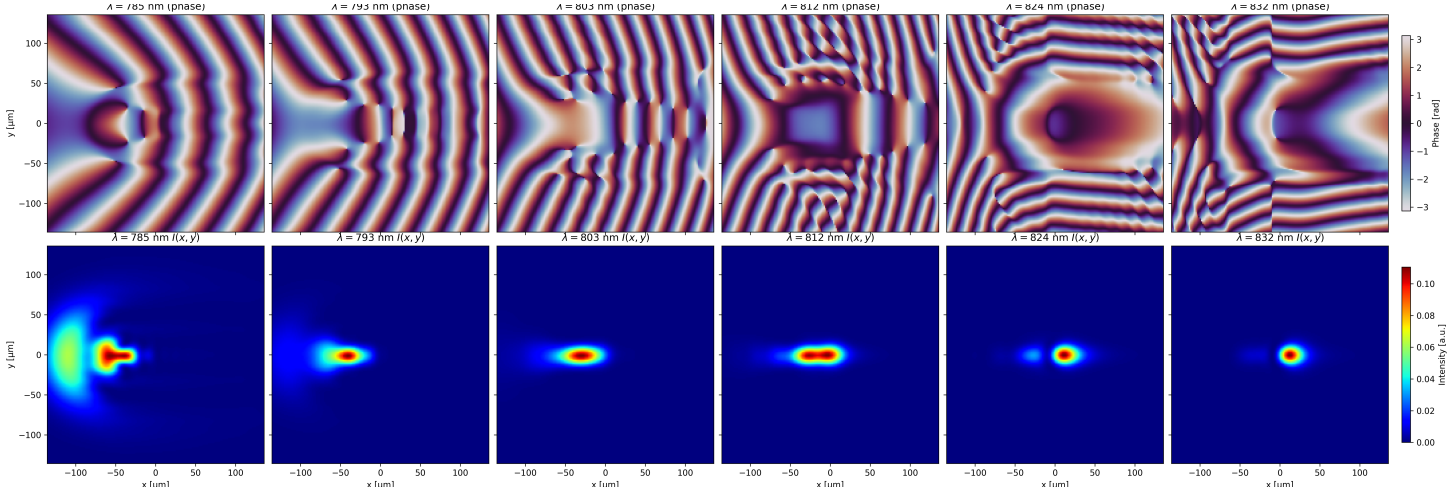

Figure 5.16: Simulation: Far-field spatial wrapped phase and intensity for different wavelengths after 1 minute of heating.

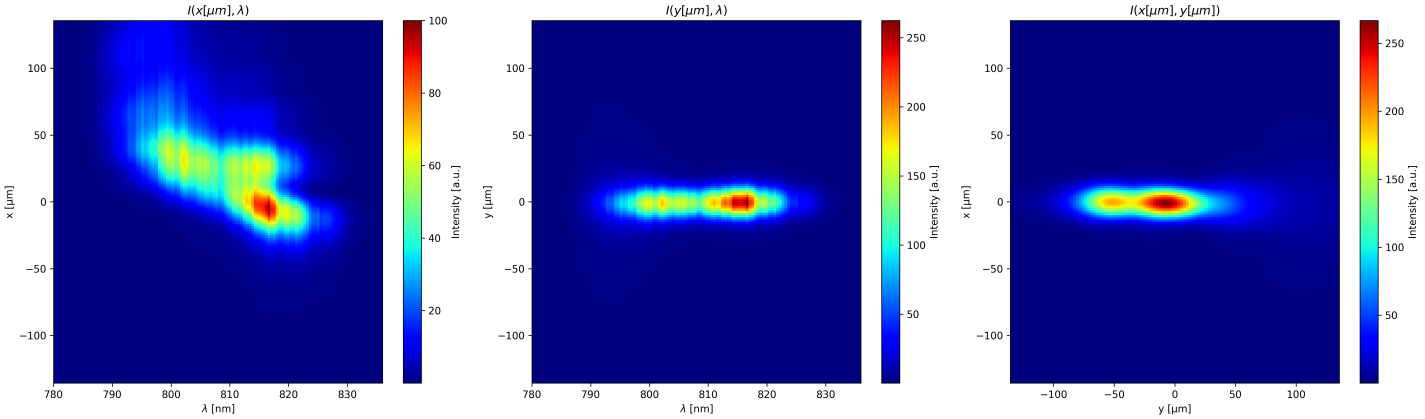

Figure 5.17: Simulation: Spatio-spectral intensity (left) and spatial intensity (right) after 1 minute of heating.

tilt dispersion, which in turn generates pulse-front tilt, as discussed in the theoretical chapter.

In conclusion, heating-induced deformation strongly impacts the laser beam by generating amplitude and phase spatio-spectral couplings, which ultimately affect the quality of the electron beam produced by the laser-plasma injector.

An important observation is that the spatio-spectral couplings are primarily caused by the second grating, since different wavelengths experience different amplitudes and derivatives of deformation on the second grating. In contrast, the first grating causes a global spatial aberration that is nearly identical across all wavelengths.

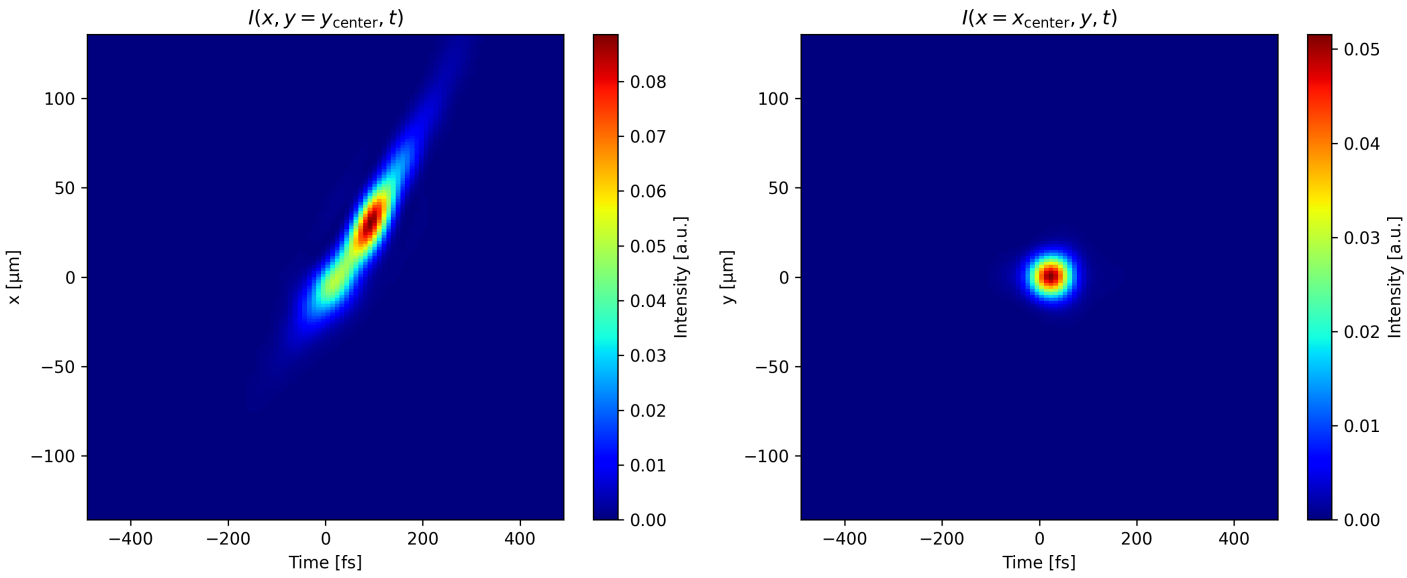

Figure 5.18: Simulation: Spatio-temporal intensity in both directions in the far field after 1 minute of heating.

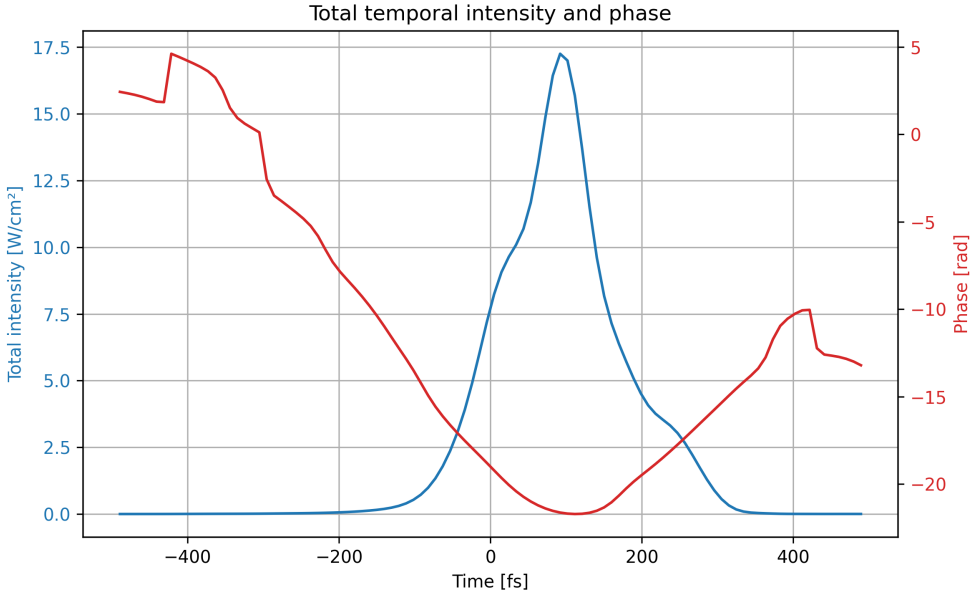

Figure 5.19: Simulation: Temporal intensity and phase after 1 minute of heating.

Finally, one may wonder why the same Zernike analysis was not performed on the measurements, even though we have access to the reconstructed phase. This was not

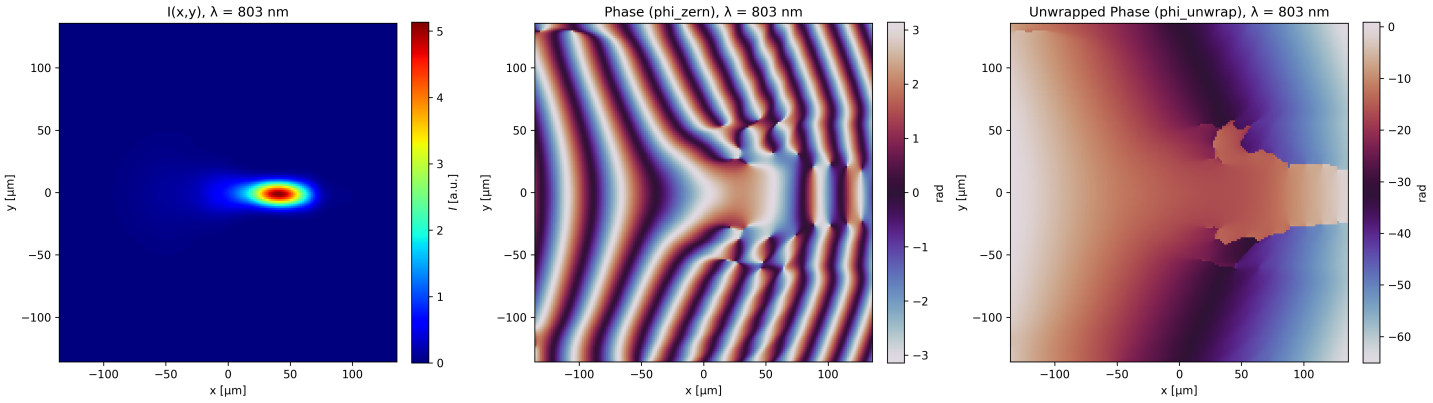

Figure 5.20: Simulation: Amplitude and phase (wrapped and unwrapped) for $\lambda = 803\ \mathrm{nm}$ after 1 minute of heating.

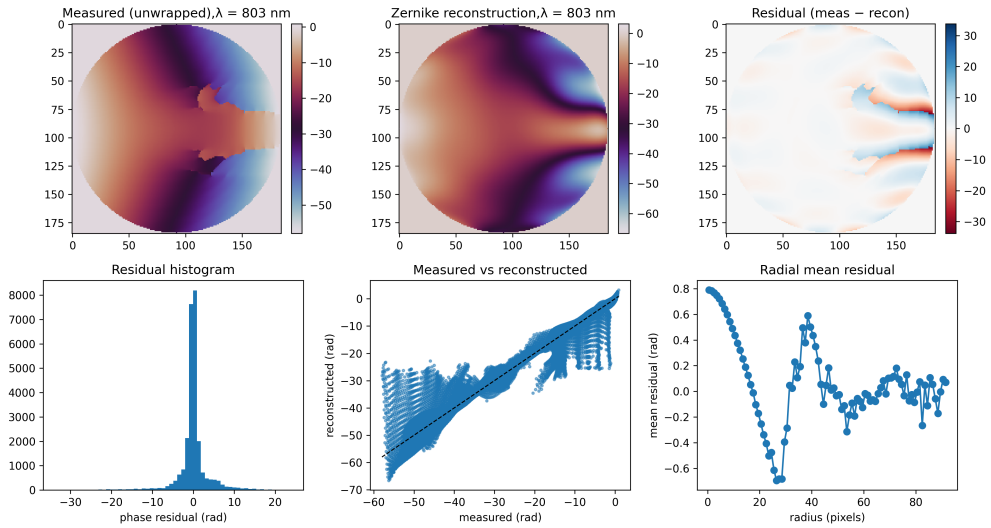

Figure 5.21: Simulation: Zernike reconstruction of the unwrapped phase of Fig. 5.20 and analysis of the residuals. The reconstruction yields an R^2 score of 0.92.

done because the phase retrieved from the GSA is not sufficiently accurate, as explained earlier. Moreover, the measured phase is so complex that achieving a good regression requires hundreds or even thousands of Zernike polynomials.

Additionally, the 1 Hz laser already exhibits aberrations. Ideally, to perform this analysis, we would need a 1 Hz laser with a very high Strehl ratio ($SR \gtrsim 0.8$). In our case, both

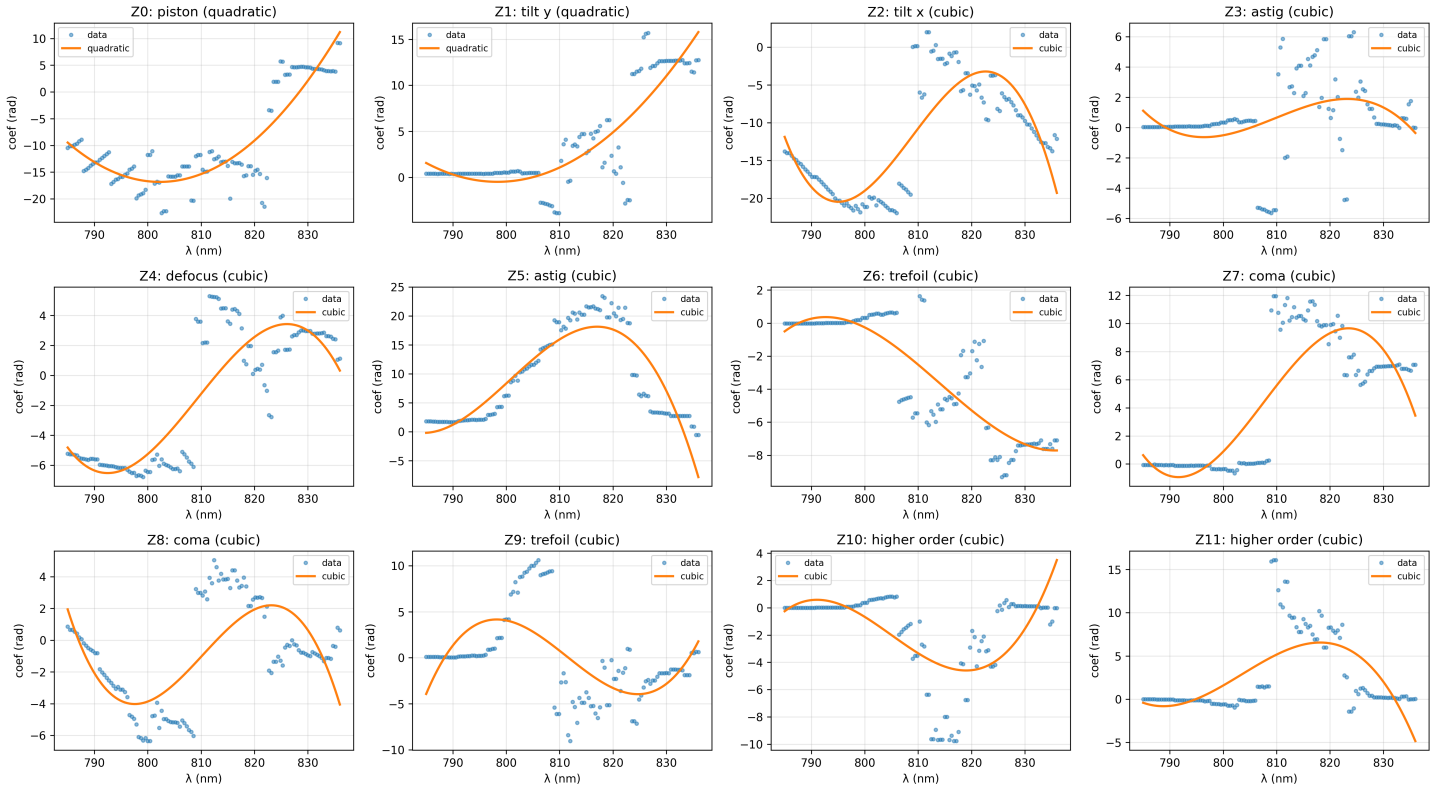

Figure 5.22: Simulation: First 12 Zernike polynomials resulting from the regression for every wavelength. Blue points correspond to individual wavelengths; the orange curve is a third-order polynomial fit.

the 1 Hz and 10 Hz measurements exhibit variations in the Zernike coefficients between wavelengths, which prevents a reliable comparison.

5.4.4 Conclusion

In conclusion, we developed a thermo-mechanical pipeline that enables the prediction of the spatio-spectral aberrations to which a laser beam is subjected as a result of thermal deformation of the compressor gratings. The code produces reliable results that are in good agreement with experimental observations. The analysis presented here, together with the simulations discussed in the next chapter, indicates that performing laser-plasma acceleration at 10 Hz or higher repetition rates requires a careful choice of the compressor substrate material. In particular, selecting a material with a low coefficient of thermal expansion is essential. Otherwise, strong thermally induced aberrations will significantly reduce the achievable focal intensity, thereby preventing efficient electron injection and acceleration in the LWFA regime.

Another interesting perspective would be to intentionally control the generated spatio-temporal couplings (STC). For instance, one could envision a setup combining compressor gratings with spatial light modulators, together with a precise STC diagnostic. The U-Net architecture presented in the previous chapter would be particularly suitable for this purpose, as it enables near–real-time phase retrieval and could therefore be implemented in a closed-loop feedback control system. Such a configuration, potentially involving multiple MSCCD devices (at least two), would allow either correction or deliberate introduction of specific STC. For example, controlled angular dispersion could be added to steer the electron beam.

Finally, the analysis presented here and in the previous chapter shows that by using at least two MSCCD placed in different planes around the focus, ideally with an increased number of spectral channels (e.g., 30), in conjunction with a spectral phase measurement device such as a FROG, it would be possible to fully reconstruct the complex electric field in single-shot and near–real time, without requiring additional diagnostic devices.

Appendix A: Effect of grating pitch variations

In addition to the height deformation $z_1(x, y)$ considered in the Modified Grating Equation (MGE), one could in principle include variations of the grating pitch d under thermal stress, i.e., $d \to d + \epsilon(x, y)$. This was initially considered; however, quite unexpectedly, even a pitch variation on the order of $10^{-10}\,\mathrm{m}$ (see Fig. 5.23) completely destroys the far-field pattern. A simple numerical test shows that such a tiny variation already dominates the phase.

Similarly to the height deformation, $\epsilon(x, y)$ is obtained from COMSOL; it represents the horizontal displacement of the Pyrex substrate. The map is rescaled according to

$$\epsilon(x, y) \equiv \epsilon(x, y)\frac{d}{dx},$$

where dx is the pixel size in meters and $d = 5.5 \times 10^{-7}\,\mathrm{m}$ is the nominal grating pitch. This rescaling accounts for the fact that each COMSOL pixel contains thousands of grooves, and we assume that all grooves within a pixel experience the same additional displacement.

Initially, I was surprised that such a small deformation could completely destroy the far-field pattern. My first hypothesis was that random pitch errors in real gratings are likely larger than the thermal-induced variations,

$$\epsilon_{\mathrm{rand}}(x, y) \gg \epsilon_{\mathrm{thermal}}(x, y),$$

and would therefore dominate, effectively making the thermal deformation negligible. My second assumption was that even if the random pitch errors were large (on the order of $10^{-9}\,\mathrm{m}$), their impact on the far-field pattern would "wash out" when integrated over many grooves, still yielding acceptable diffraction peaks.

However, a careful analysis of the integral in Eq. (5.7) for $z_1(x) = 0$ shows that this reasoning is incorrect. For clarity, let us denote the random pitch error simply by

$$\epsilon(x) \equiv \epsilon_{\mathrm{rand}}(x).$$

We reconsider the field for order $m = 1$:

$$E(\beta) \propto \int_{-Nd/2}^{Nd/2} \exp\left[-ik(\sin\alpha + \sin\beta)\,x + i\frac{2\pi x}{d + \epsilon(x)}\right] dx. \tag{5.21}$$

Since $\epsilon(x) \ll d$, we expand to first order:

$$E(\beta) \propto \int_{-Nd/2}^{Nd/2} \exp\left[-ik(\sin\alpha + \sin\beta)\,x + i\frac{2\pi x}{d}\left(1 - \frac{\epsilon(x)}{d}\right)\right] dx. \tag{5.22}$$

At the diffraction peak, where the grating equation is satisfied, the integral becomes

$$E(\beta) \propto \int_{-Nd/2}^{Nd/2} \exp\left[i\frac{2\pi x\,\epsilon(x)}{d^2}\right] dx. \tag{5.23}$$

Assume that $\epsilon(x)$ is a Gaussian random variable with zero mean, $\epsilon \sim \mathcal{N}(0,\sigma^2)$. Then

$$\mathbb{E}[e^{ik\epsilon}] = e^{-k^2\sigma^2/2}, \qquad k = \frac{2\pi x}{d^2}.$$

Therefore,

$$\mathbb{E}[E(\beta)] \propto \int_{-Nd/2}^{Nd/2} e^{-\frac{1}{2}k^2(x)\sigma^2} dx. \tag{5.24}$$

With the change of variables

$$u = \sqrt{\frac{k^2\sigma^2}{2}}, \qquad \alpha = \frac{\pi\sigma N}{\sqrt{2}\,d},$$

the integral becomes

$$\mathbb{E}[E(\beta)] \propto \frac{d^2}{\sqrt{2}\pi\sigma} \int_{-\alpha}^{\alpha} e^{-u^2}\, du. \tag{5.25}$$

Thus,

$$\mathbb{E}[E(\beta)] \propto \frac{d^2}{2\sqrt{2\pi}\,\sigma} \left[\operatorname{erf}(\alpha) - \operatorname{erf}(-\alpha)\right]. \tag{5.26}$$

For $N = 10^5$, $\sigma = 10^{-9}\,\mathrm{m}$, and $d = 5.5 \times 10^{-7}\,\mathrm{m}$ and noting that for independent gaussian error $Var[E(\beta)] = 0$ we find:

$$\frac{\mathbb{E}[E(\beta)]}{E(\beta)} \simeq 2.1 \times 10^{-3},$$

corresponding to a reduction of the diffraction peak amplitude by more than a factor of 10^3. Here, $E(\beta)$ denotes the nominal (undeformed) diffraction amplitude.

Therefore, the acceptable pitch error must be extremely small,

$$\epsilon(x,y)_{\text{acceptable}} \lesssim 10^{-12}\,\mathrm{m}.$$

This order of magnitude is consistent with the experimental results reported by Kochetkov *et al.* [11], who measured a root-mean-square non-equidistance of

$$\xi_{\text{rms}} = 7 \times 10^{-7},$$

defined as

$$\xi(x,y) = \frac{\Delta N(x,y)}{N},$$

where $N = 1/d$ is the nominal line density and $\Delta N(x, y)$ is the difference between the local and the nominal line density.

For a grating reported in their article with $N = 1200\,\mathrm{mm}^{-1}$, the nominal pitch is $d = 1/N = 833\,\mathrm{nm}$.

this corresponds to a random pitch variation of

$$\epsilon_{\mathrm{rand}} = 5.8 \times 10^{-13}\,\mathrm{m},$$

i.e., sub-picometer fluctuations, which is fully consistent with the above estimate of the acceptable pitch error. An alternative explanation could be that the spatial resolution of our COMSOL deformation maps is too coarse to resolve variations at the scale of individual grooves. Including these horizontal displacements as pitch variations then produces unphysical scattering in the simulation. Therefore, in practice, only the vertical deformation is included in the MGE, which captures the dominant diffraction effects.

I am not fully certain about this explanation, but it seems consistent with the observed numerical behavior and is the most plausible reasoning at this stage.

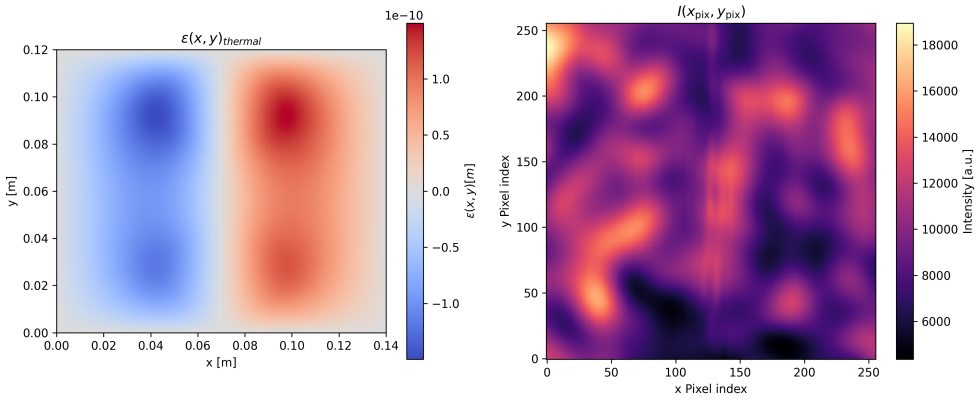

Figure 5.23: Left: Rescaled $\epsilon(x, y)$ map. Right: Resulting scattered far-field intensity pattern.

Appendix B: Solution of the modified grating equation (MGE)

Let us rearrange the MGE as:

$$\sin\beta + \frac{dz_1(x)}{dx}\cos\beta = -\sin\alpha + \frac{dz_1(x)}{dx}\cos\alpha + \frac{m\lambda}{d}. \tag{5.27}$$

Using the trigonometric identity

$$\sin(\beta + \arctan(p)) = \frac{\sin\beta + p\cos\beta}{\sqrt{1+p^2}},$$

we define

$$p = \frac{dz_1(x)}{dx}.$$

Then, the solution for β is

$$\beta = \arcsin\left(\frac{-\sin\alpha + p\cos\alpha + \frac{m\lambda}{d}}{\sqrt{1+p^2}}\right) - \arctan(p). \tag{5.28}$$

Referring to Fig. 5.6, we have $\tan\theta = \frac{\partial z}{\partial x}$. Therefore, the local diffraction angle measured from the deformed surface normal is

$$\beta_{\text{loc}} = \beta + \arctan\left(\frac{\partial z}{\partial x}\right). \tag{5.29}$$

References

[1] Daniel J. Kane and Rick Trebino. “Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating”. In: *IEEE Journal of Quantum Electronics* 29.2 (1993), pp. 571–579.

[2] Costantino Iaconis and Ian A. Walmsley. “Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses”. In: *Optics Letters* 23.10 (1998), pp. 792–794.

[3] Spencer W. Jolly, Olivier Gobert, and Fabien Quéré. “Spatio-temporal characterization of ultrashort laser beams: a tutorial”. In: *Journal of Optics* 22.10 (2020), p. 103501. doi: 10.1088/2040-8978/abad08.

[4] Erik Löfquist et al. “Correlated x-ray and electron beam steering by pulse-front tilt in laser wakefield acceleration”. In: *Physical Review Accelerators and Beams* 28.10 (2025), p. 102802.

[5] Maxence Thévenet et al. “Pulse front tilt steering in laser plasma accelerators”. In: *Physical Review Accelerators and Beams* 22.07 (2019), p. 071301. doi: 10.1103/PhysRevAccelBeams.22.071301.

[6] Daniel E. Mittelberger et al. “Laser and electron deflection from transverse asymmetries in laser-plasma accelerators”. In: *Physical Review E* 100 (2019), p. 063208. doi: 10.1103/PhysRevE.100.063208.

[7] Antoine Jeandet et al. “Survey of spatio-temporal couplings throughout high-power ultrashort lasers”. In: *Optics Express* 30.3 (2022), pp. 3262–3288. doi: 10.1364/OE.444564.

[8] Vincent Leroux et al. “Wavefront degradation of a 200 TW laser from heat-induced deformation of in-vacuum compressor gratings”. In: *Optics Express* 26.10 (2018), pp. 13061–13071. doi: 10.1364/OE.26.013061.

[9] V. Leroux, T. Eichner, and A. R. Maier. “Description of spatio-temporal couplings from heat-induced compressor grating deformation”. In: *Optics Express* 28.6 (2020), pp. 8257–8265.

[10] Max Born and Emil Wolf. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. 7th. Cambridge, UK: Cambridge University Press, 1999.

[11] Anton Kochetkov et al. “Precise characterization of diffraction grating groove patterns”. In: *Optics Express* 33.6 (2025), pp. 13673–13681. doi: 10.1364/OE.551097.

6 -Realistic PIC simulation for LWFA

6.1 Introduction

In the previous section, we presented a Machine Learning (ML) model trained on low-fidelity Particle-in-Cell (PIC) simulations. In this chapter, we focus on realistic PIC simulations based on experimental laser measurements. The objectives are twofold: first, to compare these realistic simulations with the low-fidelity ones to validate the relevance of our ML approach; and second, to investigate higher-order effects arising from laser beam asymmetries.

This chapter presents three distinct PIC simulations. The first simulation represents a low-fidelity case derived from a random scan of input parameters comprising 600 data points. This simulation updates the laser and plasma parameters of the PALLAS platform, incorporating revised laser configurations and plasma density profiles obtained from updated computational fluid dynamics (CFD) simulations. The second simulation utilises the spatial intensity and phase profiles of the laser measured at 1 Hz, in conjunction with the temporal intensity and phase profiles acquired via FROG measurements. This configuration corresponds to the scenario detailed in the chapter addressing the compressor deformation. The third simulation employs the spatio-spectral intensity and phase profiles measured at 10 Hz, aligning with the case also presented in the deformed compressor chapter.

The simulations and experimental campaigns were conducted in an unstructured manner. Unfortunately, the data collected using the laser diagnostics bench of the LIF module, following the laser energy upgrade, were obtained with a laser exhibiting significant spatial phase modulation prior to the compressor. Consequently, this study is limited to a comparative analysis using plasma parameters initially optimized for the expected laser conditions during the commissioning phase, together with a low-repetition-rate (1 Hz) reference laser of suboptimal beam quality.

The chapter is structured as follows. We begin by reviewing the basic theory of PIC simulations with azimuthal mode decomposition. Next, we present the input parameters and results for the low-fidelity simulations, followed by the experimental inputs and results for the 1 Hz laser simulations, and then for the 10 Hz laser case. Finally, I introduce in the Appendix, to the best of our knowledge, an original effect arising from free-space propagation, which I term *Self-Interference Spatio-Spectral Phase Modulation* (SISSPM).

6.2 PIC theory

The Particle-In-Cell (PIC) method relies on self-consistently solving the Maxwell-Vlasov equations, which couple the dynamics of charged particles with the evolution of the electromagnetic fields.

The Vlasov equation for a species s (electrons, ions, etc.) reads:

$$\frac{\partial f_s(\mathbf{x},\mathbf{p},t)}{\partial t} + \frac{\mathbf{p}}{\gamma m_s}\cdot\nabla_{\mathbf{x}} f_s(\mathbf{x},\mathbf{p},t) + q_s\left(\mathbf{E}+\frac{\mathbf{p}}{\gamma m_s}\times\mathbf{B}\right)\cdot\nabla_{\mathbf{p}} f_s(\mathbf{x},\mathbf{p},t) = 0, \tag{6.1}$$

where:
- $f_s(\mathbf{x},\mathbf{p},t)$ is the distribution function of species s,
- q_s and m_s are the charge and mass of the species,
- $\mathbf{E}(\mathbf{x},t)$ and $\mathbf{B}(\mathbf{x},t)$ are the electric and magnetic fields,
- $\gamma = \sqrt{1+|\mathbf{p}|^2/(m_s^2c^2)}$ is the relativistic Lorentz factor.

The fields themselves are obtained from Maxwell's equations:

$$\nabla\cdot\mathbf{E}(\mathbf{x},t) = \frac{\rho(\mathbf{x},t)}{\varepsilon_0}, \tag{6.2}$$

$$\nabla\cdot\mathbf{B}(\mathbf{x},t) = 0, \tag{6.3}$$

$$\nabla\times\mathbf{E}(\mathbf{x},t) = -\frac{\partial\mathbf{B}(\mathbf{x},t)}{\partial t}, \tag{6.4}$$

$$\nabla\times\mathbf{B}(\mathbf{x},t) = \mu_0\mathbf{J}(\mathbf{x},t) + \mu_0\varepsilon_0\frac{\partial\mathbf{E}(\mathbf{x},t)}{\partial t}, \tag{6.5}$$

where $\rho(\mathbf{x},t)$ and $\mathbf{J}(\mathbf{x},t)$ are the charge and current densities obtained self-consistently from the particle distribution.

$$\rho(\mathbf{x},t) = \sum_s q_s\int f_s(\mathbf{x},\mathbf{p},t)\,d^3p, \tag{6.6}$$

$$\mathbf{J}(\mathbf{x},t) = \sum_s q_s\int \frac{\mathbf{p}}{\gamma m_s} f_s(\mathbf{x},\mathbf{p},t)\,d^3p. \tag{6.7}$$

Solving this system of equations exactly is, in principle, possible. In practice, however, it would require an infinite amount of computational time, since the number of charged particles in a plasma is extraordinarily large.

The Particle-In-Cell (PIC) method overcomes this limitation by introducing a key approximation: the continuous distribution function is discretised into a finite number of so-called *macro-particles*. Each macro-particle represents an aggregate of a large number (typically millions) of real particles that share the same charge-to-mass ratio and follow the same equations of motion.

Within this framework, the continuous distribution function $f(\mathbf{x}, \mathbf{p}, t)$ is approximated as a finite sum over macroparticles:

$$f(\mathbf{x}, \mathbf{p}, t) = \sum_p w_p \, S_x(\mathbf{x} - \mathbf{x}_p(t)) \; \delta(\mathbf{p} - \mathbf{p}_p(t)) \,, \tag{6.8}$$

where w_p is the macroparticle weight, S_x the spatial shape function, $\mathbf{x}_p(t)$ and $\mathbf{p}_p(t)$ the position and momentum of macroparticle p, and δ the Dirac delta function in momentum space.

The macroparticles evolve under the action of the self-consistent electromagnetic fields, while their charge and current are deposited onto a spatial grid to provide the sources of Maxwell's equations. This approach drastically reduces the computational cost while preserving the essential kinetic physics of the plasma, enabling a self-consistent solution of the Maxwell–Vlasov system.

In Smilei, the algorithm advances the system in time through a repetitive loop. After the initialisation stage (at $t = 0$), all macroparticles are loaded into the simulation domain, and the electromagnetic fields are initialised on the entire computational grid. The code then iterates over successive time steps, each of which consists of the following operations:

- interpolation of the electromagnetic fields from the grid to the macroparticle positions,
- update of the macroparticle momenta and positions according to the Lorentz force,
- deposition of the resulting charge and current densities onto the grid,
- solution of Maxwell's equations on the grid to obtain the updated electromagnetic fields.

In this chapter, we employ a quasi-3D simulation based on an azimuthal mode decomposition. This approach consists in representing the fields as a truncated Fourier series in the azimuthal angle, defined as

$$F(r, \theta, t) = \sum_{m=-(N-1)}^{N-1} \tilde{F}_m(r, t) \, e^{-im\theta}, \tag{6.9}$$

where $F(r, \theta, t)$ denotes the complex field to be represented, $N = m_{\max}$ is the truncation order (i.e. the maximum azimuthal mode retained in the expansion), and $\tilde{F}_m$ are the azimuthal mode coefficients, given by

$$\tilde{F}_m(r, t) = \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} F(r, \theta, t) \, e^{im\theta} \, \mathrm{d}\theta. \tag{6.10}$$

This ansatz relies on the fact that, in many laser–wakefield acceleration (LWFA) configurations, the problem exhibits a high degree of cylindrical symmetry. As a result, the field dynamics can be accurately captured using only the first few azimuthal modes, allowing the series to be truncated at low order without significant loss of information. This drastically reduces the computational cost compared to a full 3D simulation.

Within this framework, all electromagnetic fields are expressed using the azimuthal mode decomposition, while the macroparticles still evolve in a fully six-dimensional phase space. Further details on quasi-3D PIC simulations and their implementation in Smilei can be found in [1, 2]

6.3 Gaussian temporal envelope simulation

The main goal was to compare an ideal case with two realistic cases: the first realistic case corresponds to the laser operating at 1 Hz with high-energy laser pulse characterisation data, while the second corresponds to the laser operating at 10 Hz with high energy, with a degraded laser pulse from heated compressor gratings.

6.3.1 Gaussian temporal envelope and plasma parameters

Between the commissioning campaign in July 2025 and the measurement presented in Chapter 5 in December 2025, the last amplifier (AMP3) of the LASERIX system was upgraded. This upgrade was expected to increase the laser intensity. After the compressor, the laser energy rose from approximately 0.8 J to 1.25 J, which in principle leads to a higher a_0.

In addition, during the July campaign, where we achieved the first PALLAS electrons (chapter 2), we discovered that the laser waist was smaller than previously expected: $w_0 = 11$ μm instead of the originally anticipated $w_0 = 20$ μm. Additionally, measurements from the wavefront sensor revealed that the Strehl ratio varied between $SR \in [0.3, 0.7]$, depending on the optimisation of the deformable mirror.

Consequenlty, to determine the relevant simulation parameters for this comparison, we performed a new random scan of 600 points, accounting for the updated laser and plasma parameters. Thus, compared to the previous low-fidelity simulations presented in the chapter 3, the updated laser parameters were a waist of $w_0 = 11$ μm and a maximum a_0 in vacuum ranging from $a_0 \in [1.5, 2.25]$. All other parameters remained unchanged: we still used a flat Gaussian beam of order 5 with a single azimuthal mode, one particle per cell, and the envelope approximation. The remaining variable parameters were

$$\{x_{\text{off}} \in [-1.5, 2.5]\ \text{mm}, \quad c_{N2} \in [5, 50]\%, \quad p_1 \in [20, 100]\ \text{mbar}\}\,.$$

Out of the 600 simulations in the random scan, we selected the best one, defined as the simulation achieving the highest score according to the function of merit

$$f = \frac{\sqrt{Q}E_{\mathrm{med}}}{\delta E_{\mathrm{mad}}}, \quad E_{\mathrm{med}} > 100\ \mathrm{MeV}.$$

The best spectrum is shown in Fig. 6.1, together with the corresponding realistic density profile obtained from CFD simulations by my colleague J. Serhal, which closer to the one of the current plasma cell.

This best configuration corresponds to the updated plasma parameters with expected increase of laser intensity. But due to some issue on the beam transport, leading to significant high frequency spatial phase modulation, the maximum achievable a_0 was actually $a_0 = 1.34$. In order to have a fair comparison we used the same density profile and c_{N2}, x_{off}, p_1 with $a_0 = 1.34$ in the envelope case to make the comparison.

6.3.2 Gaussian temporal envelope simulation results

To produce a reference case for the comparison under realistic laser conditions, the PIC envelope simulation was rerun with the reduced value $a_0 = 1.34$, while keeping the plasma parameters corresponding to the best configuration obtained from the random scan (Table 6.1)

x_{off}	a_0	c_{N2}	p_1
1.3 mm	1.34	32.79%	23 mbar
E_{med}	δE_{mad}	Q	ϵ_y
40 MeV	25.8%	17.8 pC	0.85 mm.mrad

Table 6.1: Input and output parameters of the Gaussian temporal envelope simulation

In Fig. 6.2, we present the evolution of the main electron-beam parameters as a function of the longitudinal position. In addition, the wakefield structure at three longitudinal positions is shown in Fig. 6.3. We begin by discussing the evolution of the median energy, which can be approximated as

$$E_{\mathrm{med}} \sim E_x L_{\mathrm{acc}}.$$

In Fig. 6.4, we show the maximum longitudinal electric field in the simulation window as a function of the longitudinal position. This figure clearly shows that E_x follows the evolution of the laser amplitude a_0, as expected. The field reaches a maximum value of approximately $100\ \mathrm{GV/m}$ when a_0 peaks. However, the median energy reaches only $E_{\mathrm{med}} \simeq 40\ \mathrm{MeV}$, since the effective acceleration length remains relatively short. Indeed, after reaching its maximum, a_0 rapidly decreases due to diffraction.

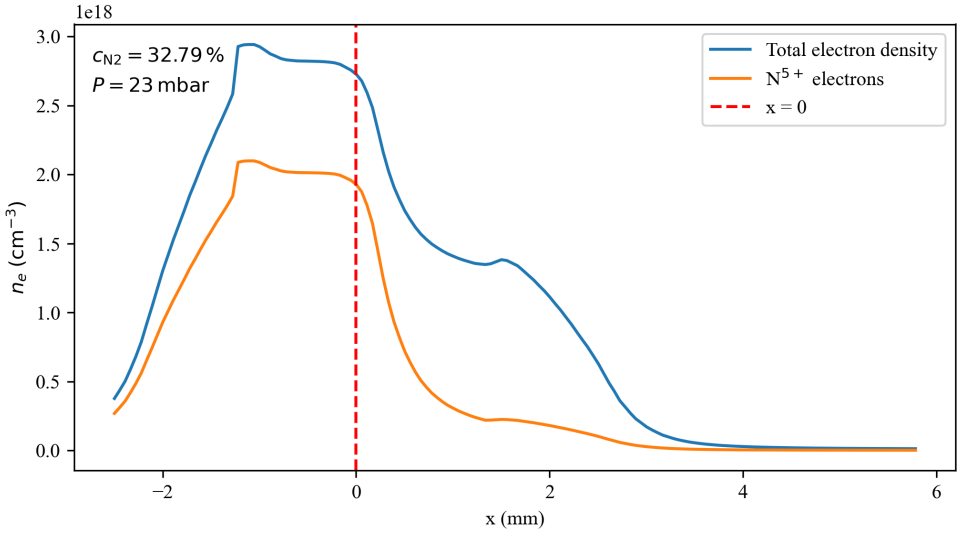

(a) Density profile from CFD simulation. The red dashed line indicates the position of the zero from which x_{off} is measured.

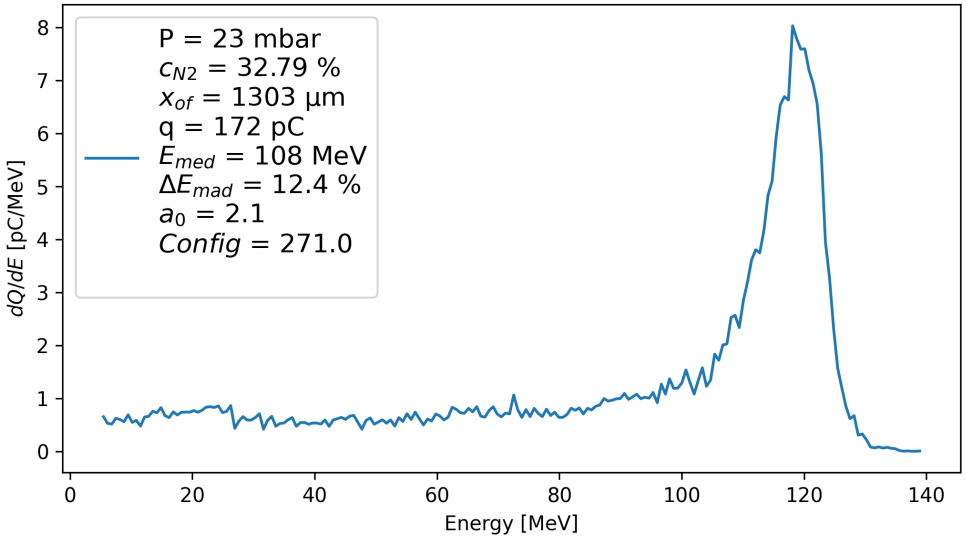

(b) Best spectrum from the 600-point random scan.

Figure 6.1: Best spectrum and density profile.

To achieve higher energies in this configuration, a_0 would need to remain high over a longer distance, for instance by extending the plateau length. When a_0 reaches its maximum, the system operates in the matched bubble regime, characterised by [3]

$$k_p w_0 \sim 2\sqrt{a_0}.$$

Using the values from Fig. 6.2 and Fig. 6.3, we find $k_p w_0 = 2.2$ and $2\sqrt{a_0} = 2.6$, indicating that the matched bubble condition is approximately satisfied. However, as the plasma density rapidly decreases downstream, this condition is no longer fulfilled and diffraction dominates, which ultimately limits the median energy to $E_{\mathrm{med}} \simeq 40\ \mathrm{MeV}$.

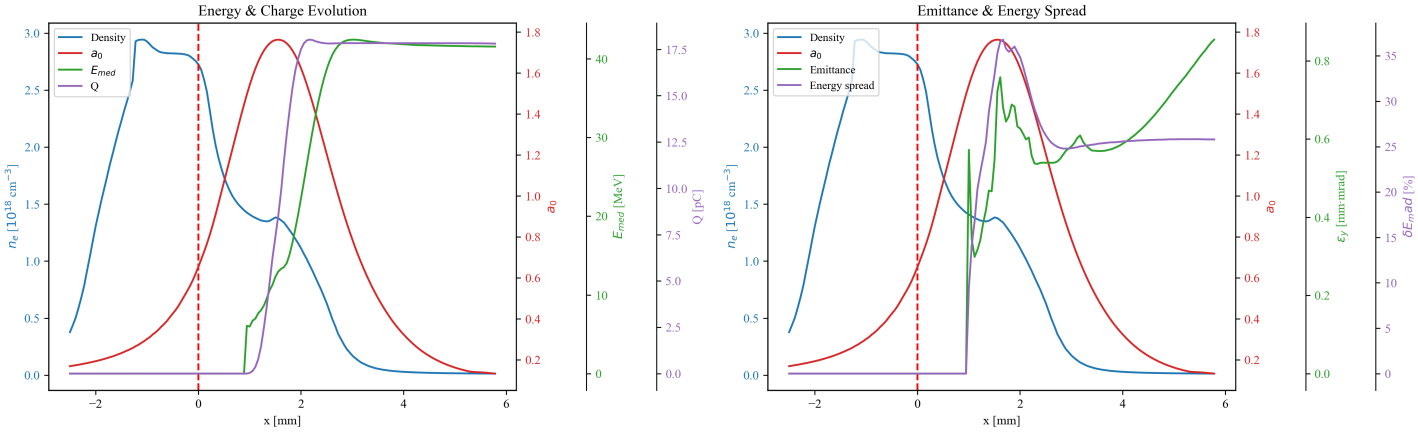

Figure 6.2: Left panel: plasma density (blue), laser amplitude a_0 (red), median energy (green), and charge (purple) for the Gaussian envelope case. Right panel: plasma density (blue), a_0 (red), emittance (green), and relative energy spread for the Gaussian envelope case.

We now turn to the evolution of the injected charge. The barrier suppression ionisation (BSI) threshold for ionising N^{6+} electrons corresponds to $a_0 = 2.21$. Since this threshold is never reached in the present simulation, all injected electrons originate from tunnelling ionisation. The tunneling ionization probability scales as $\exp(-1/E)$, and injection begins at $x = 0.94\ \mathrm{mm}$, corresponding to $a_0 = 1.5$. We therefore conclude that ionisation rates become significant only for $a_0 \gtrsim 1.5$, that is, in the region extending from $x = 0.94\ \mathrm{mm}$ to $x = 2.17\ \mathrm{mm}$, corresponding to a length of $1.22\ \mathrm{mm}$. This longitudinal interval matches precisely the region where the charge Q increases, as shown in Fig. 6.2.

Furthermore, although the tunnelling ionisation rate may remain non-negligible beyond $x = 2.17\ \mathrm{mm}$, the coefficient c_{N2} (see Fig. 6.1) becomes relatively small. As a result, electron injection would be strongly suppressed even if a_0 remains above 1.5. These combined effects explain the final injected charge of $Q = 17.8\ \mathrm{pC}$ observed in the simulation.

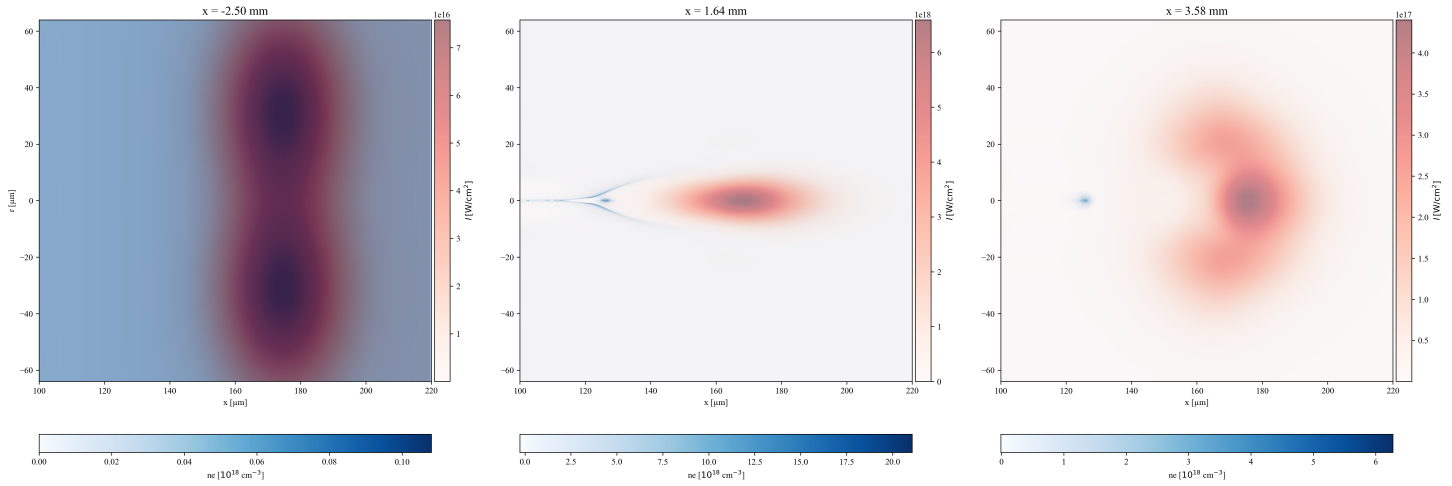

Figure 6.3: Laser intensity and wakefield structure at three longitudinal positions in the simulation for the Gaussian envelope case.

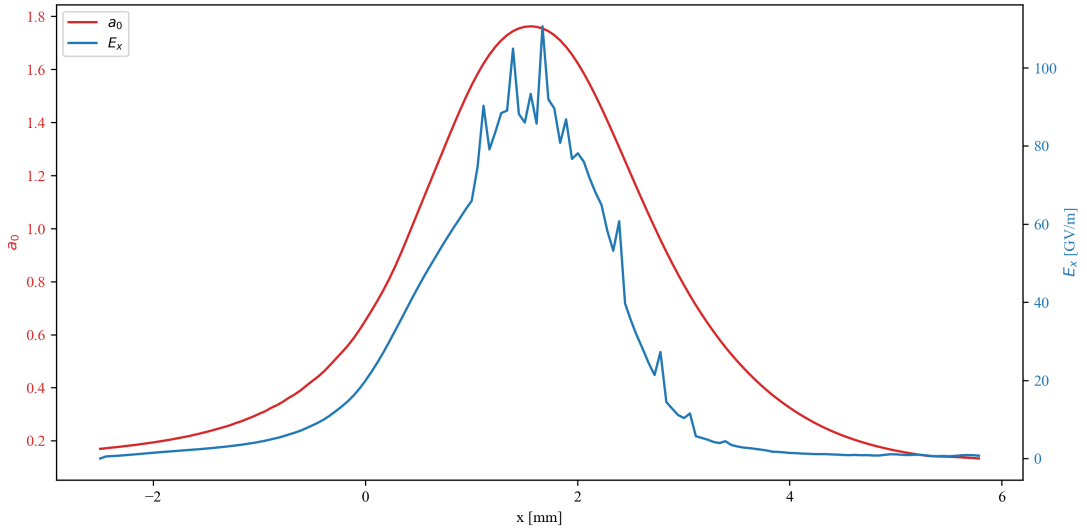

Figure 6.4: Maximum longitudinal electric field in the Smilei simulation window as a function of longitudinal position, together with the corresponding evolution of a_0, for the Gaussian envelope case.

Next, we consider the evolution of the beam emittance. Its behaviour is characteristic of electron beams produced via ionisation injection. A detailed analysis of the corresponding phase-space dynamics was provided by Xu *et al.* [4], which we briefly summarise here.

During the injection phase, the transverse emittance increases sharply and reaches a saturation value once the first electrons are injected. This rapid growth originates from the initial phase-mixing process: electrons injected at different times start their betatron oscillations with different initial phases. As a result, the occupied phase-space area increases, leading to a rise in emittance.

Subsequently, the emittance slowly decreases. Electrons injected earlier are accelerated over a longer distance and thus reach higher energies, which reduces their betatron frequency. This effect leads to a partial contraction of the transverse phase-space area. At later times, however, even electrons born at the same instant may experience variations in the accelerating field E_x. These variations induce additional differences in betatron phase, resulting in longitudinal phase mixing. As this process develops, the phase-space area gradually increases again, causing a renewed growth of the emittance.

Finally, in the present case, once acceleration ceases, space-charge forces become the dominant mechanism affecting the beam dynamics. These forces lead to a further increase of the emittance, which reaches a final value of $\epsilon_y = 0.85\ \mathrm{mm} \cdot \mathrm{mrad}$.

We now turn to the evolution of the energy spread. As observed for the emittance, the energy spread initially increases rapidly as the longitudinal phase space is populated by electrons born at different positions and/or times. These electrons experience different values of the accelerating field E_x and therefore reach different energies. After the injection ends, the beam continues to be accelerated, which results in a reduction of the *relative* energy spread as the median energy increases. Once acceleration is complete, the relative energy spread stabilises at a constant value of $\delta E_{\mathrm{mad}} = 25.8\%$. Overall, this configuration produces a typical ionisation-injection electron beam. For this plasma configuration, corresponding to the optimal set of parameters obtained from the random scan and using the experimentally available suboptimal laser intensity, a straightforward way to improve the beam quality would be to increase the length of the plasma density plateau. Doing so would extend the acceleration distance, leading to a higher median energy E_{med}, a modest increase in the injected charge Q, and a reduction of the relative energy spread δE_{mad}, all other parameters being equal.

6.4 1 Hz Laser simulation

6.4.1 1 Hz Laser parameters

The 1 Hz laser measurements were performed on the PALLAS diagnostic table using a far-field CCD camera for the spatial profile and a FROG for the temporal profile. After

the compressor grating, the laser energy was approximately $1.25\ \mathrm{J}$.

The CCD camera was scanned along the z direction over three planes ($z = -1.5, 0, +1.5\ \mathrm{mm}$, with $z = 0$ defining the focal plane). One image per plane was selected to apply the Gerchberg-Saxton Algorithm (GSA) and retrieve the spatial phase. The measured intensity and reconstructed GSA fields are shown in Fig. 6.5. While the reconstruction is not perfect, with relative squared amplitude errors of 11.3% and 5.7% for the backward and forward planes, respectively, the retrieved phase provides a physically plausible focus suitable for our simulations.

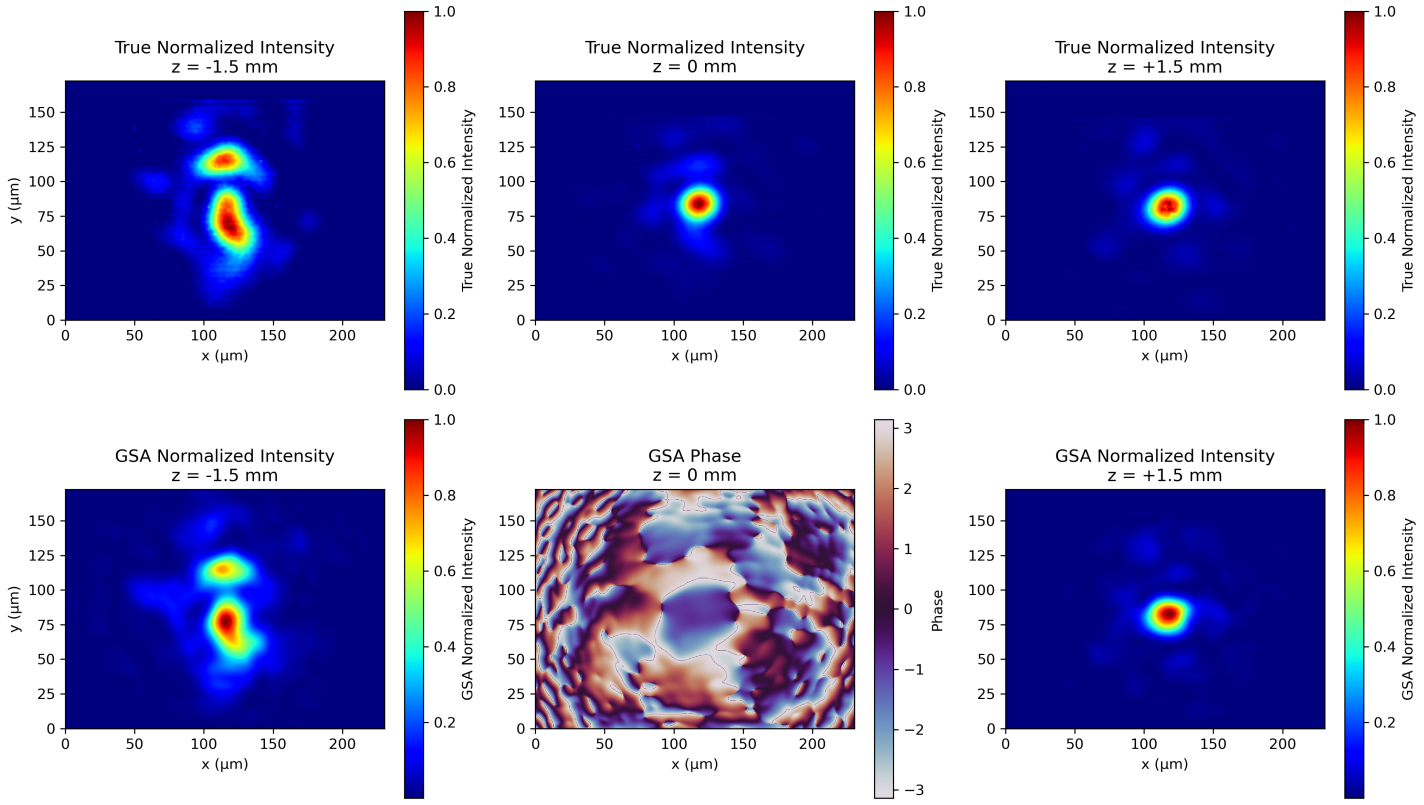

Figure 6.5: Comparison between the true normalised intensity and GSA-reconstructed fields. Top row: true normalized intensity at $z = -1.5\,\mathrm{mm}$, $z = 0\,\mathrm{mm}$, and $z = +1.5\,\mathrm{mm}$. Bottom row: GSA normalized intensity at $z = -1.5\,\mathrm{mm}$, GSA phase at $z = 0\,\mathrm{mm}$, and GSA normalized intensity at $z = +1.5\,\mathrm{mm}$.Note the resemblance between the CCD pattern and the MSCCD measurement in FIG. 5.8, which is expected since they were taken at same time and approximately the same plane.

Since $x_{\mathrm{off}} = 1.3\ \mathrm{mm}$, this corresponds to a focus located at $x = 3.78\ \mathrm{mm}$ in Fig. 6.1, because x_{off} is measured relative to the red dashed line in that figure. Accordingly, the

intensity profile was propagated backward by a total distance of $x = -3.78$ mm, as illustrated in Fig. 6.6,FIG. 6.7.

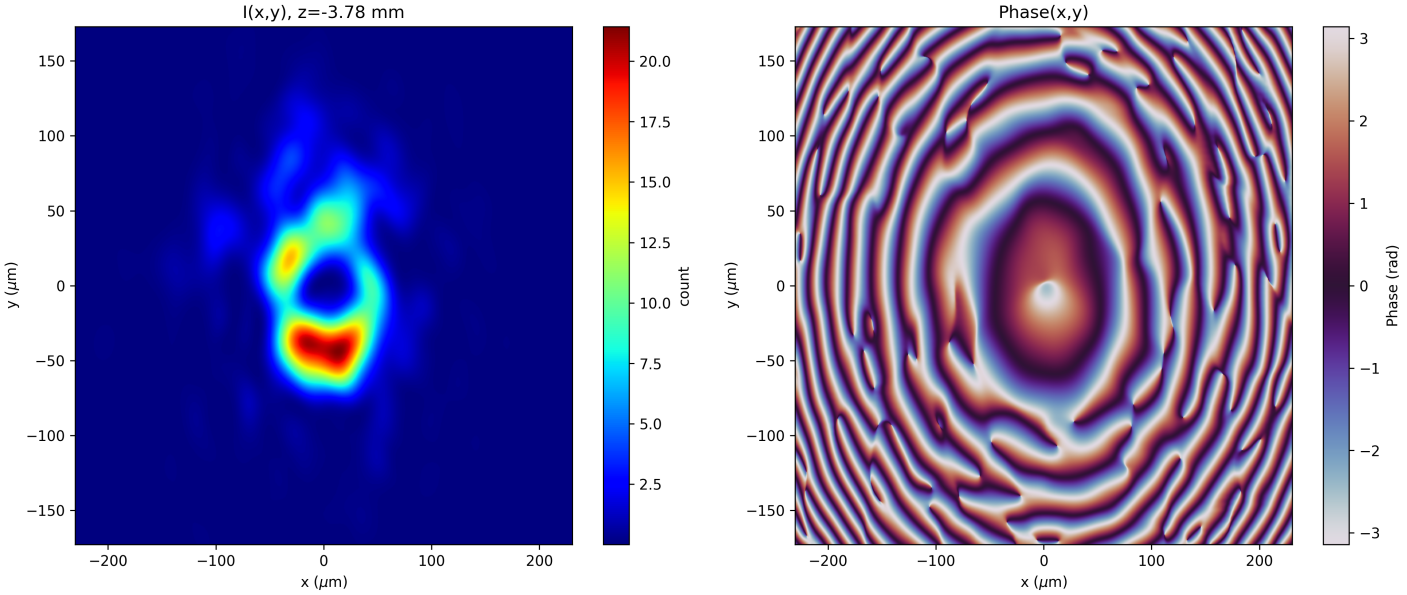

Figure 6.6: Left panel: spatial intensity after a backward propagation of -3.78 mm. Right panel: corresponding spatial phase.

The temporal intensity and phase were obtained from FROG measurements during the July campaign, representing the best achievable temporal profile for the LASERIX system (Fig. 6.8).

The spatial and temporal profiles were combined using the LASY Python library[5]. A module was added to include the spatial phase, producing the electric field

$$E(x,y,t) = \sqrt{I(x,y)}\, e^{i\phi(x,y)}\, \sqrt{I(t)}\, e^{i\phi(t)}. \tag{6.11}$$

Since no significant spatial chirp was observed (cf. Compressor Deformation chapter), this procedure assumes approximate separability of the field, which is reasonable for the 1 Hz laser.

LASY then converts the electric field from a Cartesian grid to a polar representation (r, θ, t) and performs an azimuthal mode decomposition. The laser energy is rescaled to the desired 1.25 J, corresponding to a focused vector potential of $a_0 = 1.34$.

We selected 7 azimuthal modes, as this number provides a good match between the reconstructed pulse and the original measurement, as shown in Fig. 6.9. Finally, the LASY azimuthal modes $\tilde{E}_{y,m}$ are converted into $\tilde{B}_{r,m}$ and $\tilde{B}_{\theta,m}$ for Smilei.

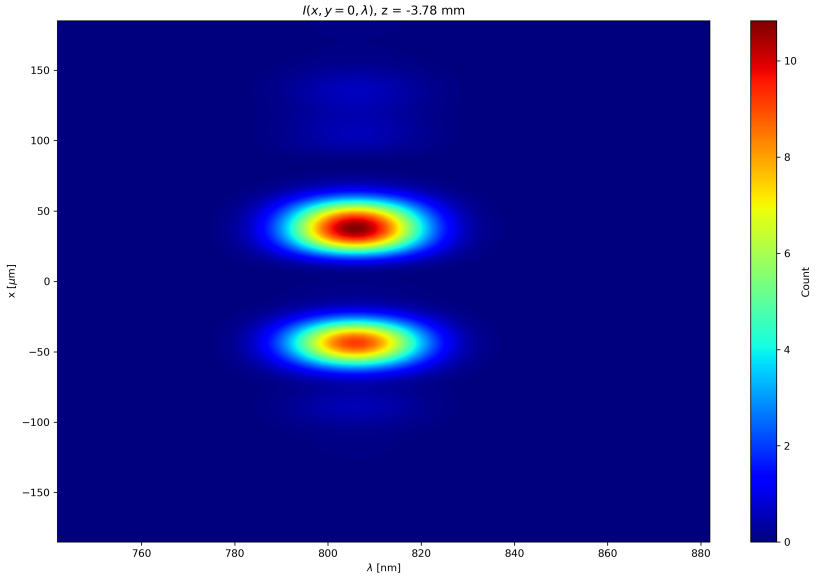

Figure 6.7: Spatio spectral intensity at the position $y = 0$ after a backward propagation of -3.78 mm. As expected there is no spatio-spectral coupling in this case

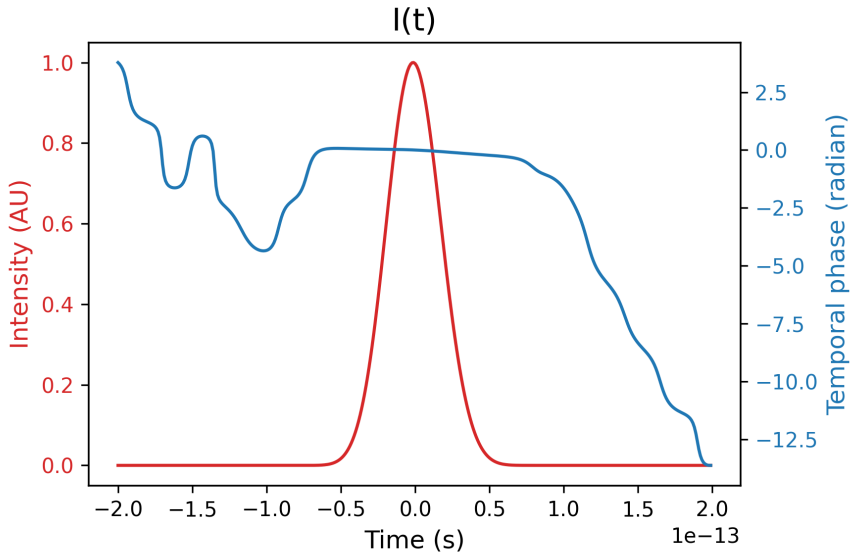

Figure 6.8: Temporal intensity and phase obtained from the FROG measurement.

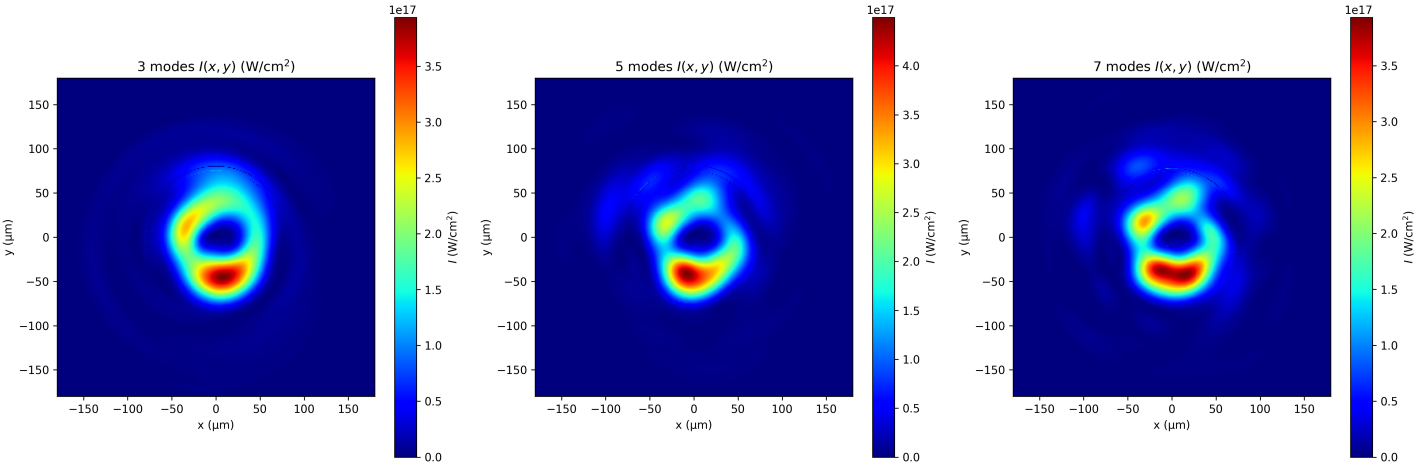

Figure 6.9: Laser reconstructed spatial intensity with different numbers of azimuthal modes: from left to right, 3, 5, and 7 modes.

6.4.2 1 Hz Laser simulation results

In Table. 6.2, we present the input and output parameters of the simulation. First, let us examine the median energy and charge for this case, as presented in Fig. 6.10. We also show the wakefield at three longitudinal positions in Fig. 6.11.

x_{off}	a_0	c_{N2}	p_1	E_{med}	δE_{mad}	Q	ϵ_y
1.3 mm	1.34	32.79%	23 mbar	30 MeV	25%	22 pC	1.5 mm.mrad

Table 6.2: Input and output parameters of the Gaussian temporal envelope simulation

In Table. 6.2, we present the input and output parameters of the simulation. First, let us examine the median energy and charge for this case, as presented in Fig. 6.10. We also show the wakefield at three longitudinal positions in Fig. 6.11.

The beam reaches a final energy and charge of 30 MeV and 22 pC, corresponding to a slightly higher charge than the ideal case, and a slightly lower energy than the ideal Gaussian envelope case. To understand the energy behavior, let us examine the maximum longitudinal electric field as a function of longitudinal position, shown in Fig. 6.12.

Focusing on the right side of the figure, the evolution is as expected: the longitudinal field grows with a_0, in agreement with 1D nonlinear theory and numerous wakefield simulations.

However, note the two large spikes in the longitudinal electric field that occur before a_0 becomes significant and before any wakefield is generated in the blowout regime.

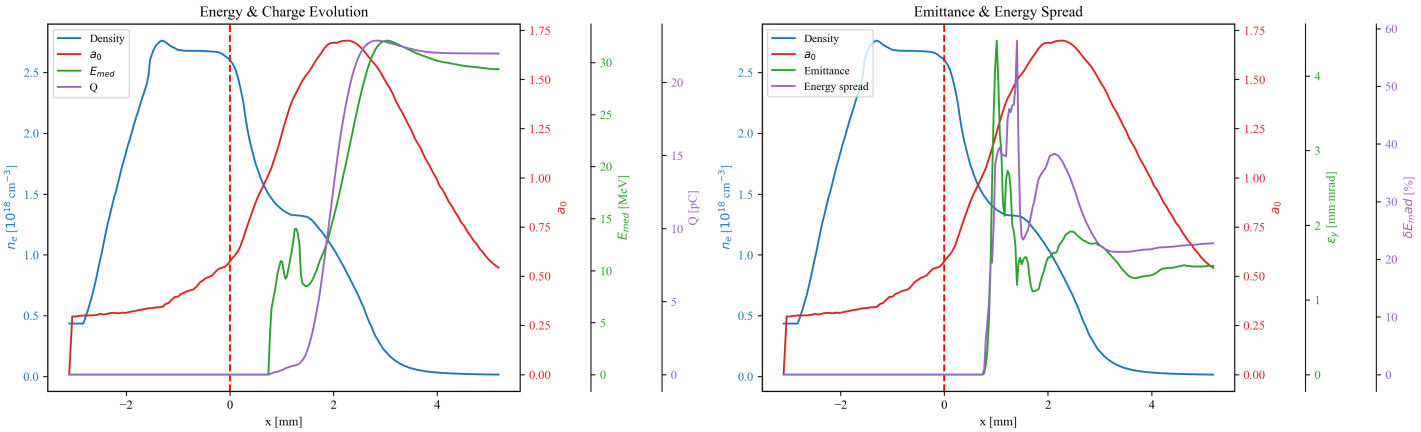

Figure 6.10: Left pannel: Density in blue, a_0 in red, median energy in green, and charge in purple for the 1 Hz case. Right panel: Density in blue, a_0 in red,emittance in green, and relative energy spread for the 1 Hz case

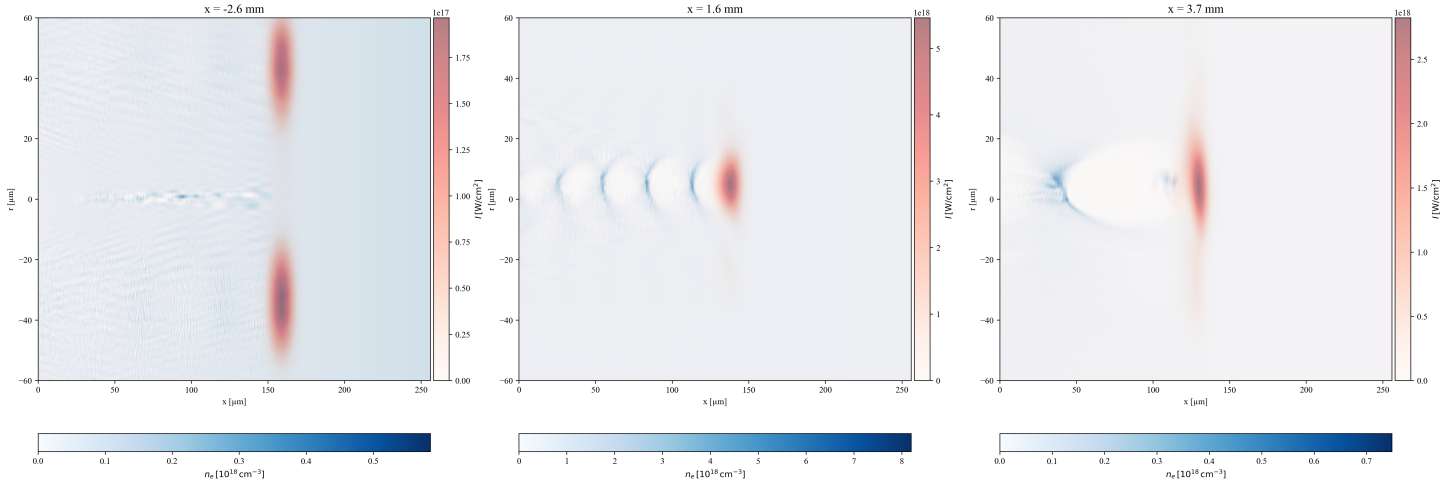

Figure 6.11: Laser Intensity and wakefield at three positions in the simulation for the 1 Hz case.

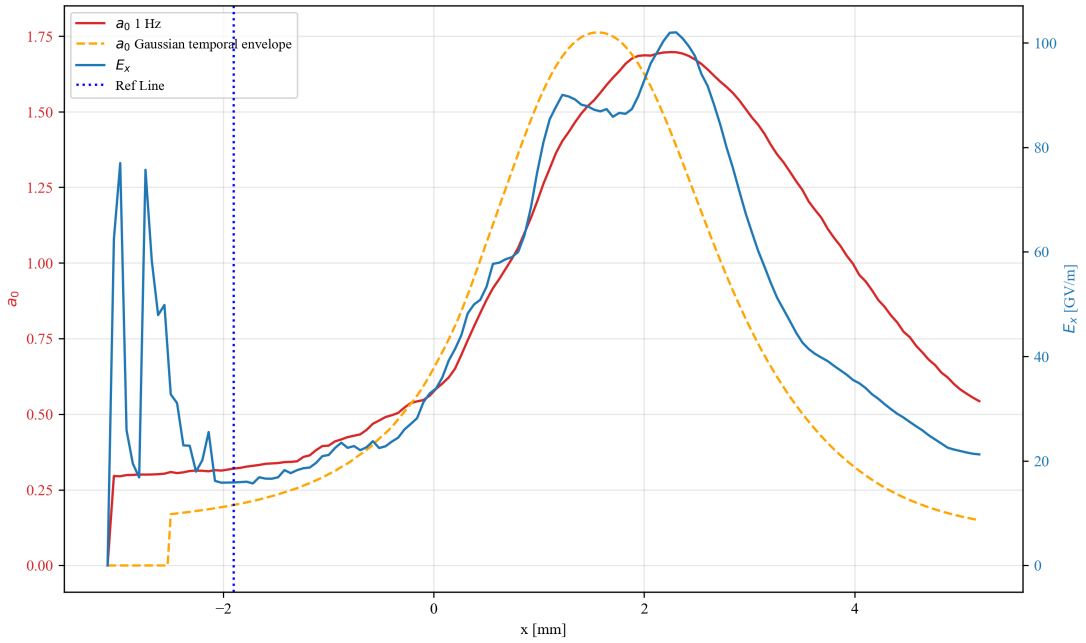

Figure 6.12: Maximum longitudinal electric field in the Smilei window for each time step of the simulation, together with a_0 for the 1 Hz case in red. The dashed orange curve corresponds to a_0 in the Gaussian temporal envelope case

This naturally raises the question: are these spikes numerical artifacts, or are they physical in origin?

In order to investigate this question, we performed a simulation in vacuum using the same laser, with a finer time step, extending up to the dashed line in Fig. 6.12. We present the results of the vacuum simulation in the Appendix C

We now discuss the evolution of E_{med}, Q, ϵ_y, and δE_{mad}. The underlying mechanisms are similar to those observed in the Gaussian envelope case discussed previously. As can be seen by comparing Fig. 6.2 and Fig. 6.10, the overall evolution of these quantities is qualitatively similar and characteristic of an ionization-injection electron beam.

Focusing on the differences, the injected charge is slightly higher in the 1 Hz case. This is because the condition $a_0 > 1.5$ is satisfied over a longer longitudinal distance, approximately $1.5\,\mathrm{mm}$, leading to a final charge of $Q = 22\,\mathrm{pC}$. The median energy, on the other hand, is slightly lower in the 1 Hz case. This reduction arises because electrons slip out of phase with the wake more rapidly than in the Gaussian envelope configuration.

In both cases, a significant fraction of the acceleration occurs in the density down-ramp. As the plasma density n_e decreases, the plasma wavelength λ_p increases, causing the electrons to progressively slip out of phase with the accelerating region of the wake. In the 1 Hz case, a larger fraction of the acceleration takes place in the down-ramp because the laser amplitude a_0 reaches its maximum at a later position (see FIG. 6.12). As a result, electrons experience a decelerating field for a longer distance. This effect explains the slight decrease of E_{med} observed in Fig. 6.10 beyond $x = 2.9\,\mathrm{mm}$, leading to a final value of $E_{\mathrm{med}} = 30\,\mathrm{MeV}$.

The evolution of the relative energy spread δE_{mad} and the emittance ϵ_y follows trends similar to those observed in the Gaussian envelope case. The initial spikes in δE_{mad} and ϵ_y can be disregarded, as they occur when the total injected charge is negligible ($Q \approx 0\,\mathrm{pC}$). The final values are $\delta E_{\mathrm{mad}} = 25\%$ and $\epsilon_y = 1.5\,\mathrm{mm}\cdot\mathrm{mrad}$.The slightly larger emittance compared to the Gaussian case can be attributed to the larger transverse size of the laser beam, which increases the injection volume. Since ionization occurs only where the laser field exceeds the ionization threshold ($a_0 \gtrsim 1.5$), a broader laser profile enlarges the transverse region where electrons can be ionized and trapped.

The most significant difference between the two configurations arises from the development of spatio-temporal couplings, in particular (SISSPM; see Appendix D), and the subsequent formation of pulse-front tilt (PFT) in the 1 Hz case. The careful reader will notice that, at $x = 3.7\,\mathrm{mm}$, the laser exhibits a clear signature of PFT in the 1 Hz configuration (see FIG. 6.11). In the following, we explain why PFT develops in this scenario.

In Fig. 6.13, we show the laser intensity above 30% of its maximum value, together with the centroid position, at three different longitudinal positions in the plasma. The figure clearly demonstrates that pulse-front tilt develops during propagation, and even reveals the onset of pulse-front curvature.

To understand why PFT develops even for a laser beam that initially exhibits no

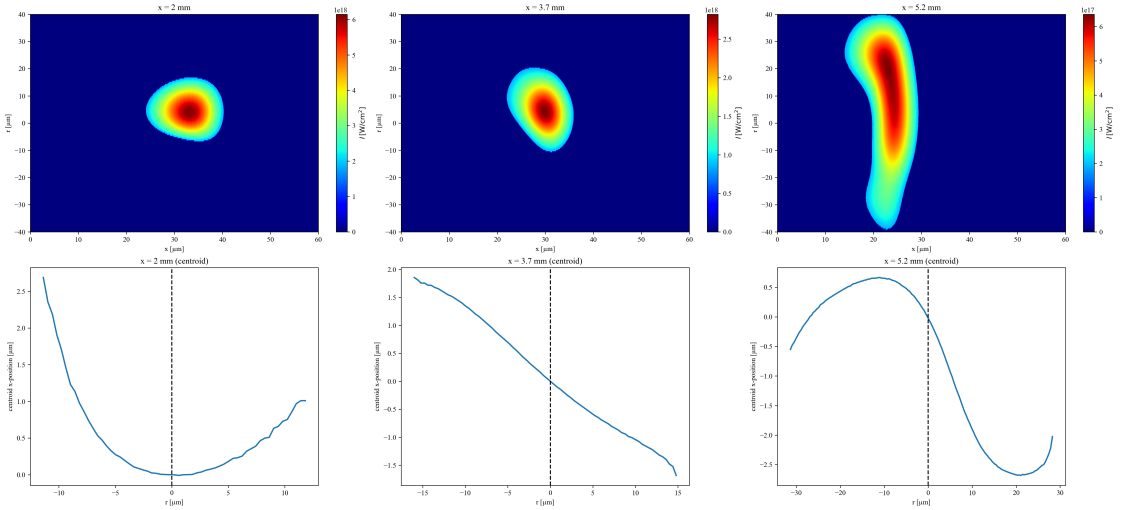

Figure 6.13: Top row: transverse laser intensity at three longitudinal positions in the plasma, with a mask showing only regions where the intensity exceeds 30% of the maximum. Bottom row: corresponding centroid positions for each transverse slice.

spatio-temporal coupling, it is useful to recall the physical origin of pulse-front tilt. Numerous studies have shown that PFT arises from angular dispersion (AD) [6, 7, 8, 9, 10], which itself develops when a laser pulse propagates through a dispersive medium characterized by a frequency-dependent refractive index, $\eta \equiv \eta(\omega)$. For example, the refractive index of a simple glass prism can be expanded to first order as

$$\eta(\omega) = \alpha(\omega - \omega_0).$$

If the derivative $\mathrm{d}\eta/\mathrm{d}\omega|_{\omega_0} = \alpha \neq 0$, the laser pulse experiences angular dispersion and consequently develops pulse-front tilt.

In a plasma, the refractive index inside the laser pulse is given by

$$\eta(\omega, r, z) \simeq 1 - \frac{\omega_p^2}{2\omega_0^2\sqrt{1 + A^2(\omega)}}\frac{n}{n_0}, \tag{6.12}$$

where ω_p is the plasma frequency and $A(\omega)$ is the normalized vector potential.

Focusing on the frequency dependence only, that is $\eta(\omega) \equiv \eta(\omega, r = \mathrm{const}, z = \mathrm{const})$, Eq. (6.12) can be written as

$$\eta(\omega) \simeq 1 - \frac{1}{\sqrt{1 + I(\omega)}}\, C(r, z), \tag{6.13}$$

where $C(r, z)$ contains the spatial dependence.

We see that the only frequency dependence of the refractive index arises from the laser pulse spectrum. In other words, if the central frequency evolves such that

$$\left.\frac{\mathrm{d}\eta}{\mathrm{d}\omega}\right|_{\omega_0} = \alpha \neq 0,$$

the pulse will develop angular dispersion and therefore pulse-front tilt (PFT).

In Fig. 6.14, we present the temporal electric field at two transverse positions, $y = 0\ \mu\mathrm{m}$ and $y = 14\ \mu\mathrm{m}$, for three different longitudinal positions in the plasma. We clearly observe here the signature of self-phase modulation (SPM): the instantaneous frequency exhibits a positive chirp, meaning that lower frequencies arrive earlier in time.

In Fig. 6.15, we show $\eta(\omega)$ at three different longitudinal positions, as defined in Eq. (6.13), with $C(r, z) = 1$ and a normalized spectrum $I(\omega)$. As the laser propagates and undergoes SPM, its central frequency shifts, leading to a progressive increase of $\mathrm{d}\eta/\mathrm{d}\omega|_{\omega_0}$. This clearly demonstrates that the frequency dependence of the refractive index becomes stronger during propagation.

Finally, Fig. 6.16 shows both the pulse-front tilt and the quantity $\mathrm{d}\eta/\mathrm{d}\omega|_{\omega_0}$ as functions of the longitudinal position. The two curves closely resemble each other, confirming the direct link between angular dispersion and PFT. In addition, $\mathrm{d}\eta/\mathrm{d}\omega|_{\omega_0}$ clearly precedes the development of PFT: it reaches its first minimum at approximately $x = 1.5\,\mathrm{mm}$, whereas the PFT reaches its first minimum around $x = 3\ \mathrm{mm}$. This delay is expected, since angular dispersion must accumulate over a finite propagation distance before manifesting as pulse-front tilt.

As discussed in numerous studies, when a laser pulse develops PFT while propagating in a plasma, the accelerated electron beam experiences transverse steering [11, 12, 13] . This leads to the final major difference between the 1 Hz and Gaussian envelope cases. In the Gaussian envelope configuration, the final mean transverse momentum is $p_y = 1 \times 10^{-3}\ \mathrm{MeV}$, whereas in the 1 Hz case it reaches $p_y = 1.4\ \mathrm{MeV}$. This results in a small but non-negligible steering of the electron beam.

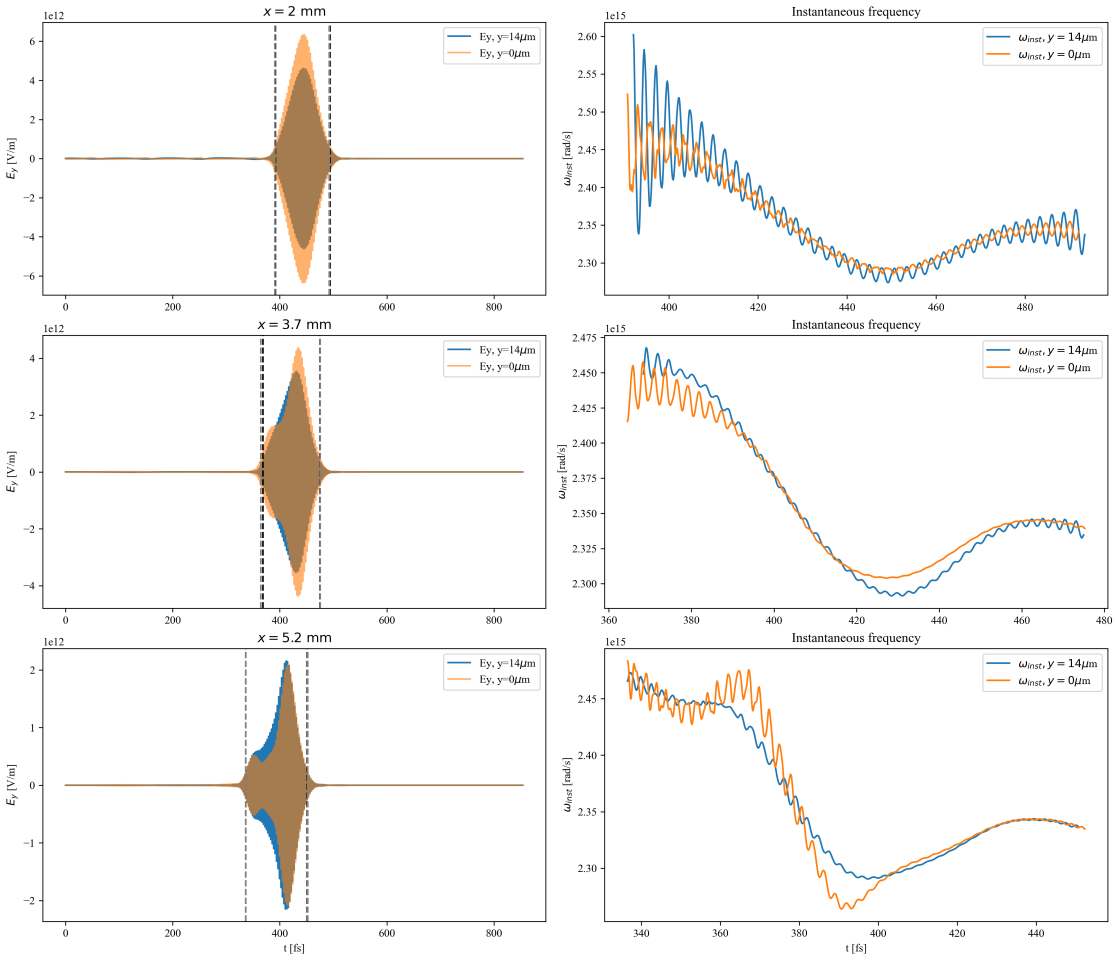

Figure 6.14: Left panel: transverse electric field at $y = 0\ \mu\mathrm{m}$ (orange) and $y = 14\ \mu\mathrm{m}$ (blue). Right panel: corresponding instantaneous frequency. Each row corresponds to a different longitudinal position in the plasma.

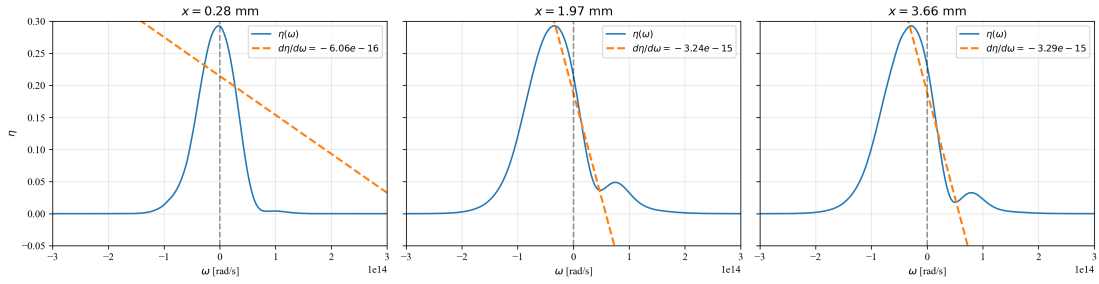

Figure 6.15: $\eta(\omega)$ (blue) as defined in Eq. (6.13). The dashed orange line represents the local slope at $\omega = \omega_0$. Each panel corresponds to a different longitudinal position.

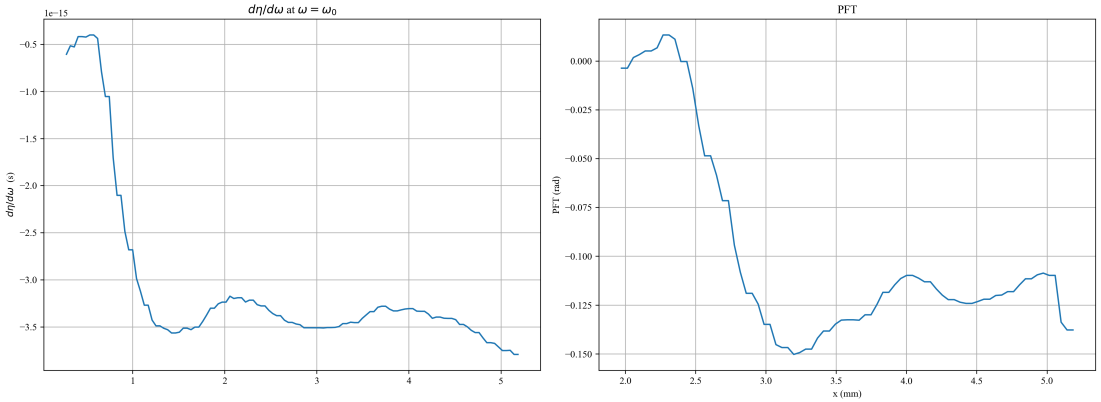

Figure 6.16: Left panel: $\mathrm{d}\eta/\mathrm{d}\omega|_{\omega_0}$ as a function of the longitudinal position. Right panel: pulse-front tilt as a function of the longitudinal position.

6.5 10 Hz Laser simulation

6.5.1 10 Hz Laser parameters

This measurement was performed after 30 minutes of compressor heating with a 10 Hz laser (cf. Compressor Deformation chapter). For this measurement, the instruments used were the multi-spectral camera and a spectrometer. We did not use the FROG, because the diagnostic table was not properly set up at the time. To generate the realistic input we used the spatio-spectral profile presented in FIG. 5.12 and the procedure presented in the previous chapter.

Smilei requires a temporal electric field. Since the GSA-retrieved phase is defined up to a global piston term, we have:

$$\phi_{\mathrm{True}}(x, y, \omega_i) = \phi_{\mathrm{GSA}}(x, y, \omega_i) + \phi(\omega_i). \tag{6.14}$$

In other words, this is a degree of freedom of the algorithm since different pistons will give the same result. After propagation the algorithm is unable to retrieve this piston term. The problem for us is that this additional term carries a lot of physical meaning, since this is the spectral phase defined as:

$$\phi(\omega_i) = \phi(x = 0, y = 0, \omega) \quad or \quad \phi(\omega) = \arg\left(\iint_{-\infty}^{+\infty} E(x, y, \omega) dx dy\right)$$

If we do not include this term, the temporal intensity will be arbitrary and depend on the piston chosen by the GSA. Unfortunately, since the FROG was not available during the compressor heating measurement, we do not directly have access to this information. The solution I adopted is to replace this piston term with the spectral phase ob-

tained from the deformed compressor code after 30 minutes of heating[1] , as presented in Fig. 6.17.

After this step, we use the measured spectrum shown in Fig. 6.17/FIG. 5.9 to scale the intensity. The field is then propagated to $z = -3.78\ \mathrm{mm}$, similarly to the 1 Hz case (see Fig. 6.18). Next, a Fourier transform is performed to obtain the temporal field.

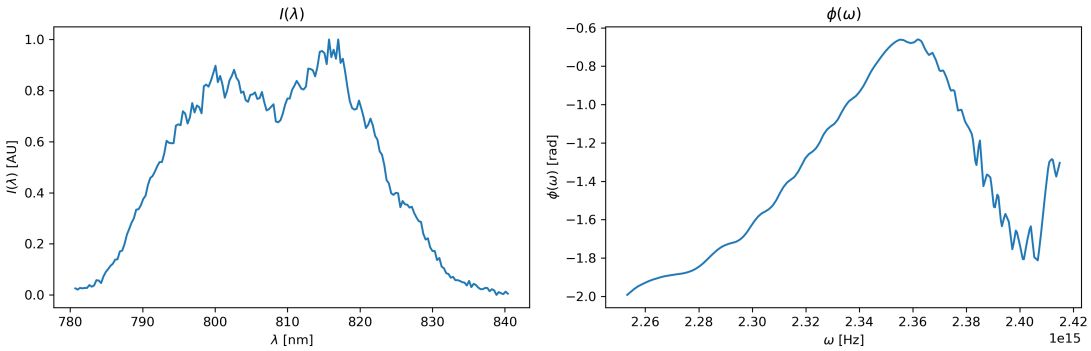

Figure 6.17: Left panel: Measured intensity spectrum. Right panel: Spectral phase from the deformed compressor code.

To achieve a high temporal resolution of $dt = 0.5\ \mathrm{fs}$ with this array, significant zero-padding is required. We used a padding factor of 64, which results in a long computation for the electric field; fortunately, this calculation was performed on a server with sufficient CPU and memory resources. The resulting temporal intensity is shown in Fig. 6.19.

The energy is scaled using the same method as for the 1 Hz case. Finally, the electric field is converted from a Cartesian grid to a polar grid, and, for the same reasons discussed in the 1 Hz case, we used 7 azimuthal modes. A three-dimensional comparison of the resulting pulse intensity with an ideal Gaussian reference is shown in Fig. 6.20.

[1]The spectral phase shown here is obtained from the deformed compressor code, but with slightly different parameters, since this chapter was written before the previous one. As a result, the phase is not quantitatively identical to the case presented in the previous chapter, although it is qualitatively similar ($\phi_2, \phi_3 \neq 0$).

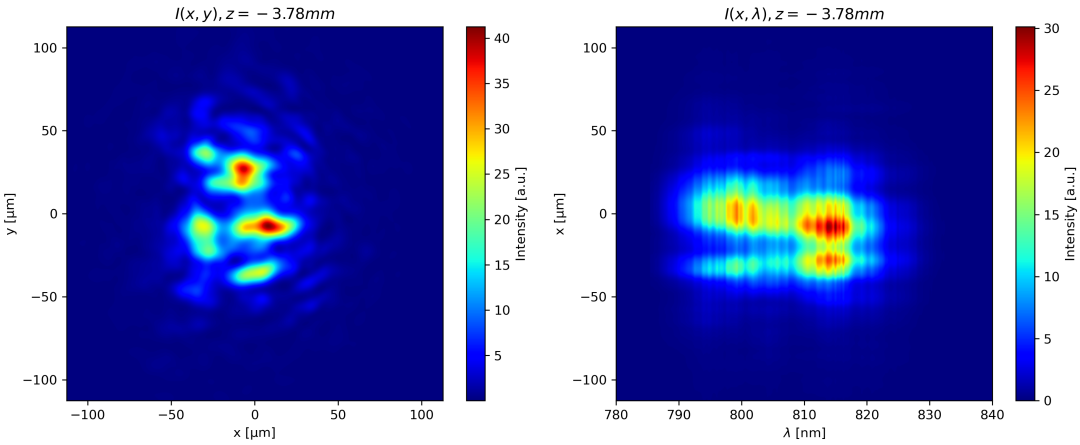

Figure 6.18: Left panel: Spatial intensity after backward propagation of $-3.78\,\mathrm{mm}$ for the multi-spectral intensity. Right panel: corresponding spatio-spectral intensity along the direction of the spatial chirp.

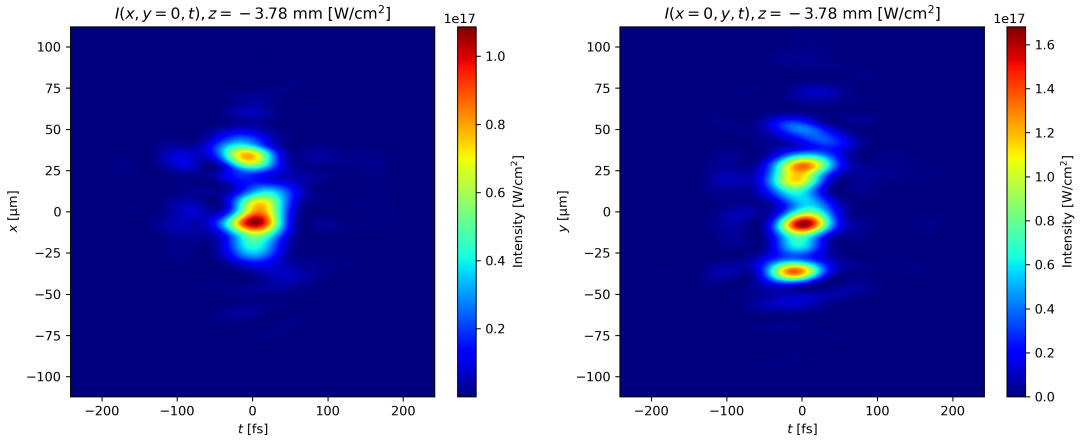

Figure 6.19: Final spatio-temporal intensity in both directions, representing the Smilei input.

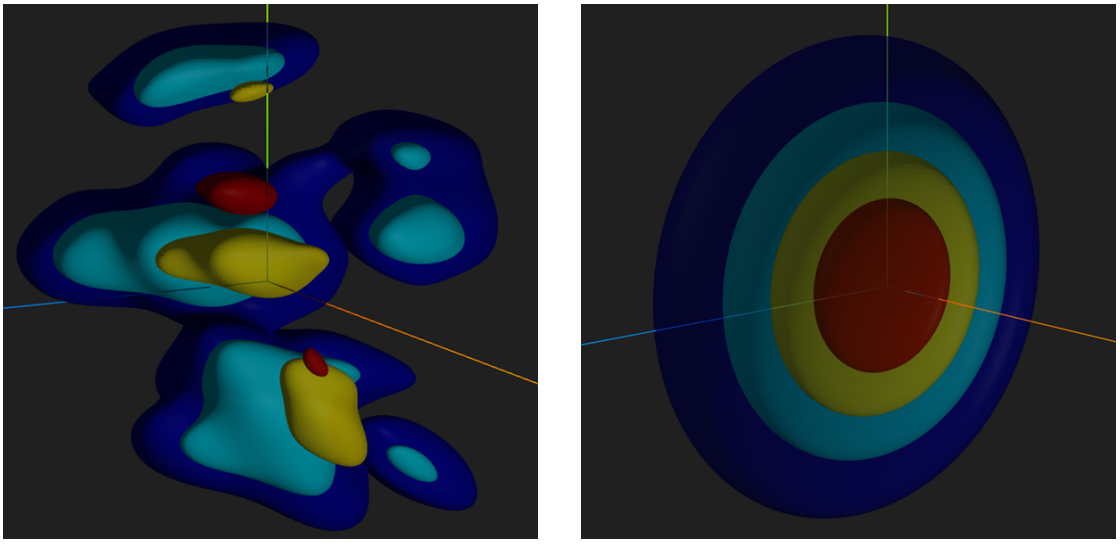

(a) Realistic aberrated laser pulse used for the 10 Hz simulation.

(b) Reference Gaussian pulse for comparison.

Figure 6.20: Three–dimensional visualization of the laser pulse intensity used in the 10 Hz PIC simulations. Left: experimentally reconstructed aberrated beam used for the 10 Hz simulation. Right: ideal Gaussian reference beam. The surfaces represent intensity isosurfaces of the normalized pulse intensity. Each color represents a different level: blue: $0.2\,|E_{\max}|^2$, cyan: $0.4\,|E_{\max}|^2$, yellow: $0.6\,|E_{\max}|^2$, red: $0.8\,|E_{\max}|^2$.

To the best of our knowledge, this is the first PIC simulation for LWFA that incorporates a fully non-separable, experimentally measured spatio-spectral electric field as the laser input.

6.5.2 10 Hz Laser simulation results

In Table. 6.3, we present the input and output parameters of the simulation. Due to limited time, we provide only a brief analysis of the $10\ \mathrm{Hz}$ case. In Fig. 6.22, we present the evolution of the main beam parameters as a function of the longitudinal position.

The first important observation is that, for a total laser energy of $1.25\ \mathrm{J}$, no electrons are injected and no blowout is generated. This can be simply explained by the fact that the strong aberrations reduce the maximum achievable a_0, preventing ionization injection. To provide further insight, we performed a simulation with the same laser parameters but with the total energy increased to $5\ \mathrm{J}$.

In this high-energy scenario, the evolution of the beam parameters differs significantly from the Gaussian envelope and the 1 Hz cases, as expected. Despite these differences, the absolute values of the parameters remain of the same order of magnitude

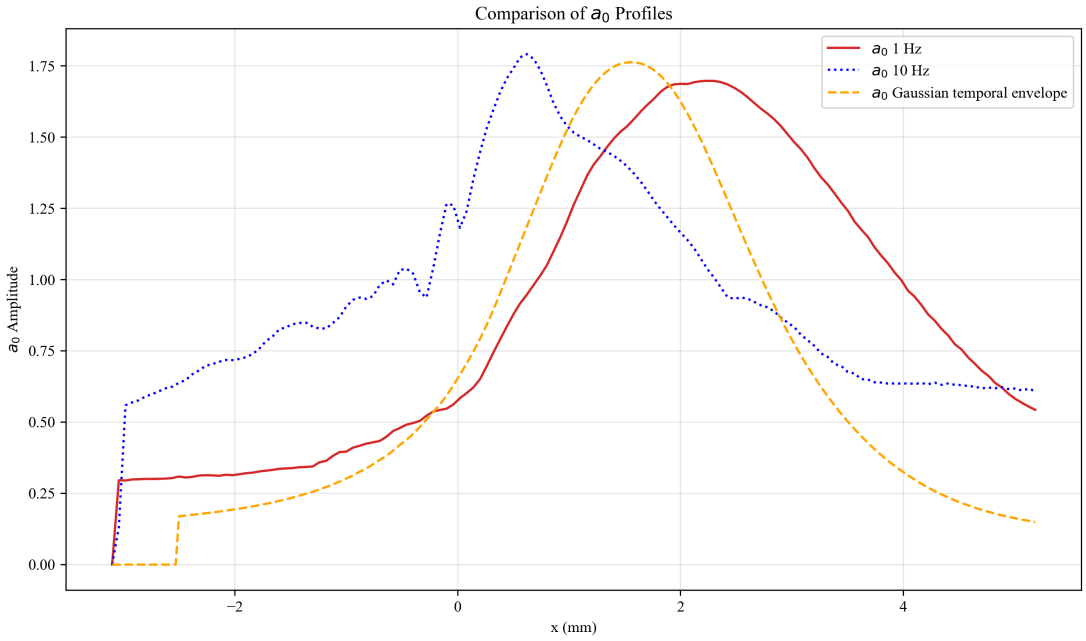

Figure 6.21: Comparison of a_0 for the three configurations at each time step of the simulation, with a_0 for the 1 Hz case shown in red. The dashed orange curve corresponds to a_0 for the Gaussian temporal envelope case. The dashed blue curve corresponds to a_0 for the 10 Hz case with a total laser energy of 5 J.

x_{off}	a_0	c_{N2}	p_1	E_{med}	δE_{mad}	Q	ϵ_y
1.3 mm	1.07	32.79%	23 mbar	32 MeV	27%	27 pC	7.2 mm.mrad

Table 6.3: Input and output parameters of the 10 Hz laser simulation with an energy of 5J

as in the previous cases. This can be explained by the fact that increasing the total laser energy to 5 J leads to a similar maximum a_0 inside the plasma, as seen in Fig. 6.21.

In Fig. 6.23, we present the laser intensity and the wakefield at three different longitudinal positions. The strongly aberrated beam initially generates three distinct wakefields. Finally, Fig. 6.24 shows the corresponding electron beam spectra for the three analyzed cases.

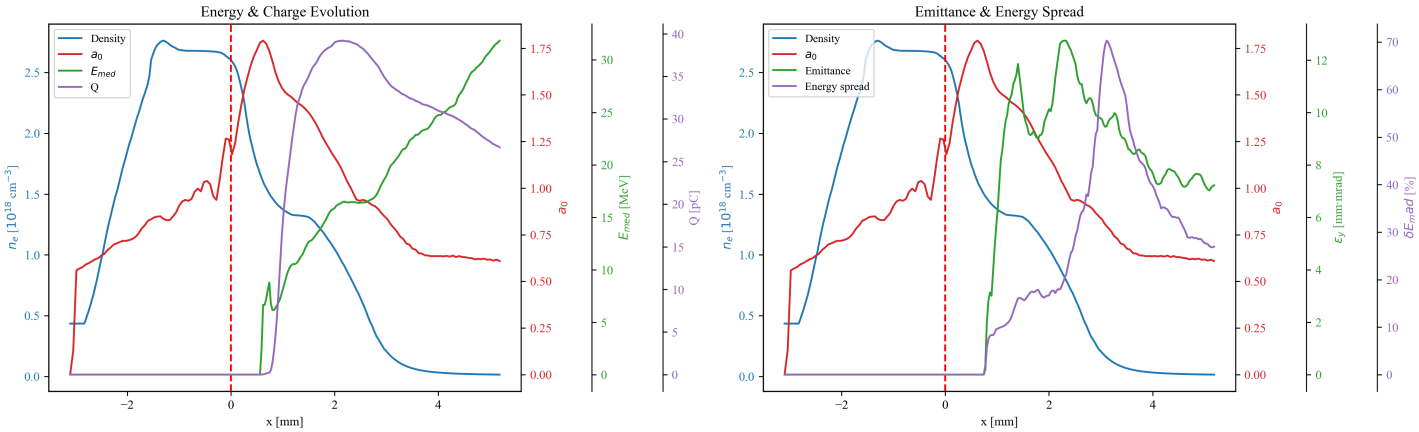

Figure 6.22: Left panel: Density (blue), a_0 (red), median energy (green), and charge (purple) for the 10 Hz case. Right panel: Density (blue), a_0 (red), emittance (green), and relative energy spread for the 10 Hz case.

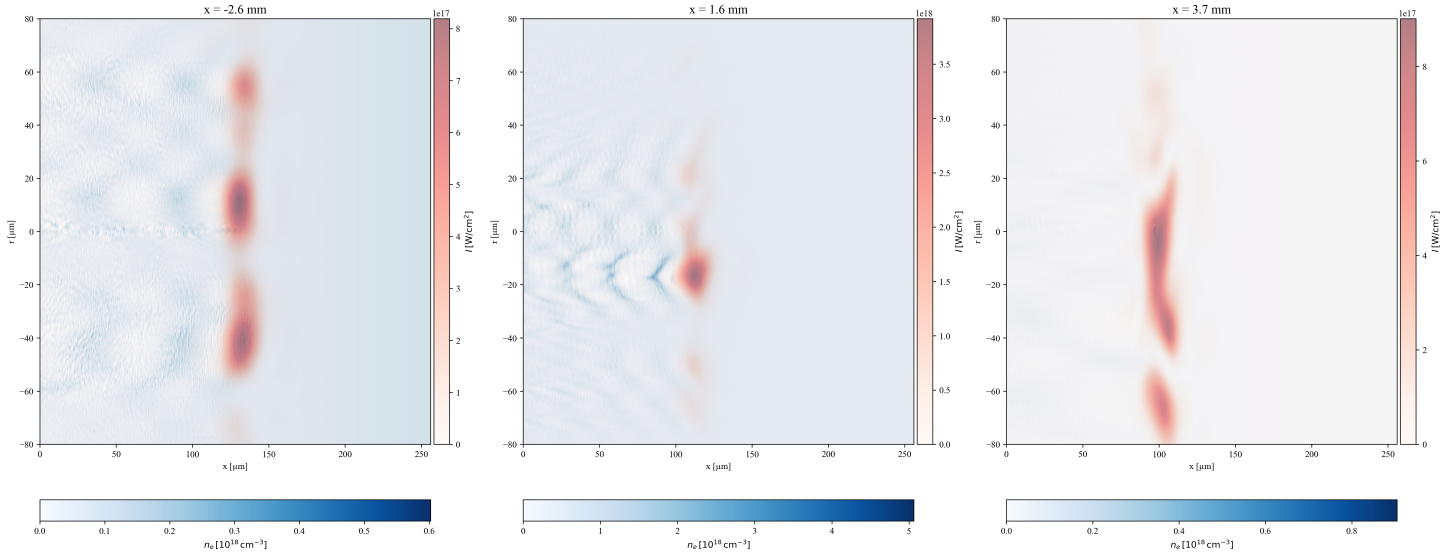

Figure 6.23: Laser intensity and wakefield at three longitudinal positions in the $10\ \mathrm{Hz}$ simulation.

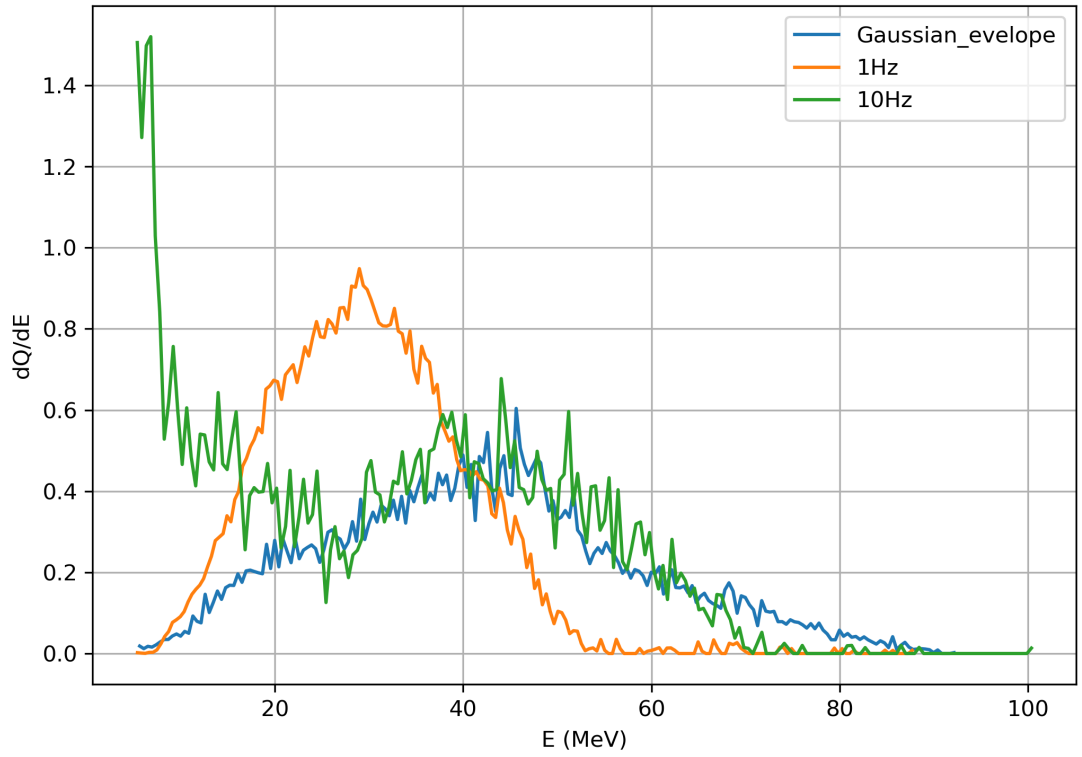

Figure 6.24: Electron beam spectra for the Gaussian envelope,1 Hz and 10 Hz case respectively.

6.6 Conclusion

In this chapter, we presented realistic PIC simulations based on experimental laser measurements and proposed a methodology to generate such realistic inputs directly from measured laser beam profiles.

The comparison between the 1 Hz realistic case and the ideal Gaussian envelope case shows that, despite some differences, the main electron beam parameters remain relatively close in the two scenarios. This observation, obtained with plasma parameters not optimized for the final laser conditions, further supports the use of simplified modeling based on the envelope approximation and strengthens the approach developed in Chapter 3.

In the 10 Hz case, the simulation showed that, for our 40 TW-class laser, electron injection and acceleration do not occur due to the strong aberrations induced by thermal deformation of the compressor, thereby confirming the results presented in the previous chapter.

Furthermore, I explained why pulse-front tilt (PFT) develops in a plasma even for a beam that initially exhibits no spatio-temporal couplings. This effect arises from self-phase modulation (SPM), which shifts the central frequency of the laser pulse, thereby inducing angular dispersion and resulting in PFT.

Finally,in the Appendix I developed a preliminary model explaining the observed self-interference spatio-spectral phase modulation (SISSPM).

Appendix C: Vacuum simulation

The first observation from the vacuum simulation is that the E_x spikes are recovered, as shown in Fig. 6.25. This indicates that the effect is independent of the plasma. Next, we examine 2D slices at the longitudinal position corresponding to the dashed red line, as presented in Fig. 6.26.

The first notable feature is that whenever the longitudinal electric field is strong, it occurs at the position of the "hole," i.e., the local minimum along the y direction. Second, at $x = -2.7\ \mathrm{mm}$, we observe that at the "hole" position the transverse electric field E_y extends far beyond the temporal duration $\tau_{\mathrm{FWHM}} = 40\ \mathrm{fs}$, indicating a strong phase modulation effect.

Focusing solely on $E_y(t, y = \mathrm{const})$, we examine the temporal electric field at two positions: $y = 0\ \mu\mathrm{m}$, corresponding to the intensity minimum, and $y = 46\ \mu\mathrm{m}$, corresponding to an intensity maximum, as shown in Fig. 6.27. Whenever the strong longitudinal electric field is present, a clear phase modulation is observed. When the longitudinal field diminishes, the phase modulation disappears. Furthermore, this effect generates a pronounced negative chirp within the pulse, with the highest frequencies arriving first.

In the following section, we present a model aiming to explain this phenomenon, which we term *self-interference spatio-spectral phase modulation (SISSPM)*.

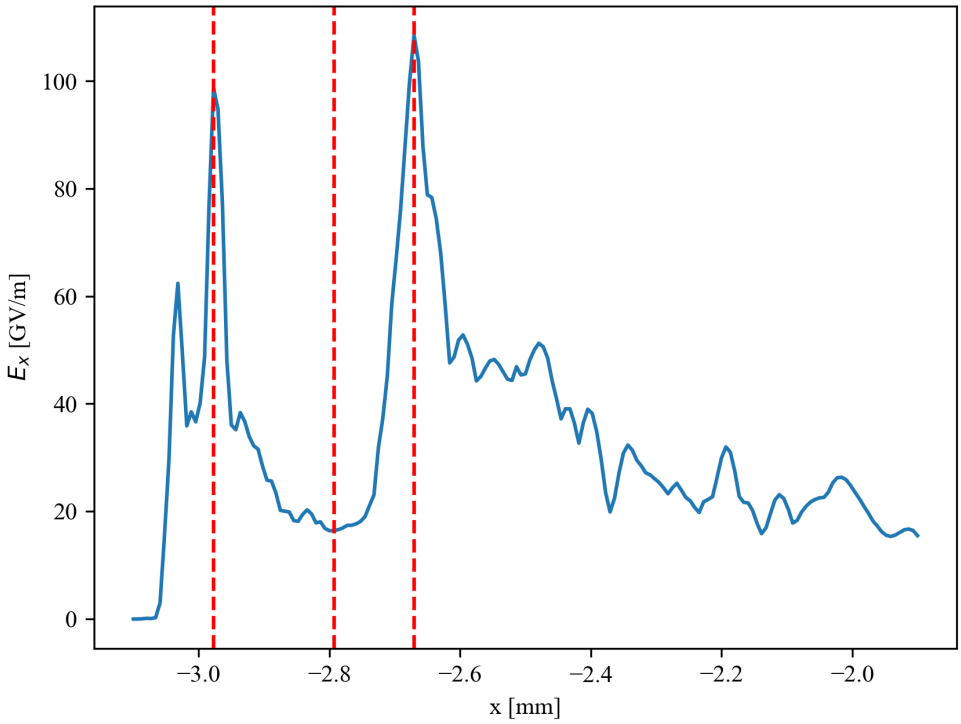

Figure 6.25: Longitudinal electric field in the vacuum simulation.

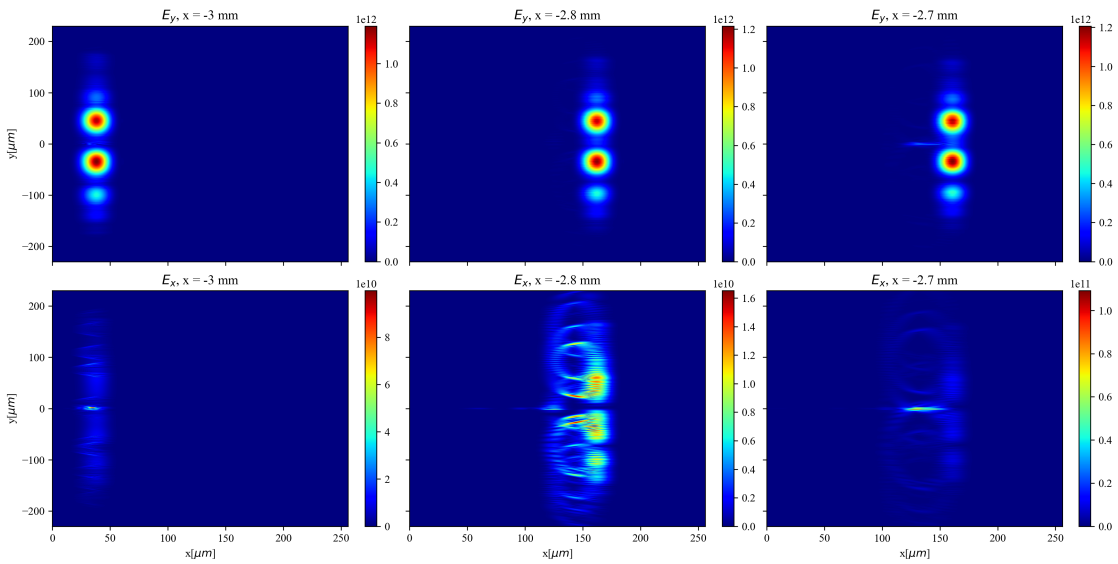

Figure 6.26: Top row: Transverse electric field (E_y) at three positions corresponding to the red dashed line in Fig. 6.25. Bottom row: Longitudinal electric field (E_x) at the same positions.

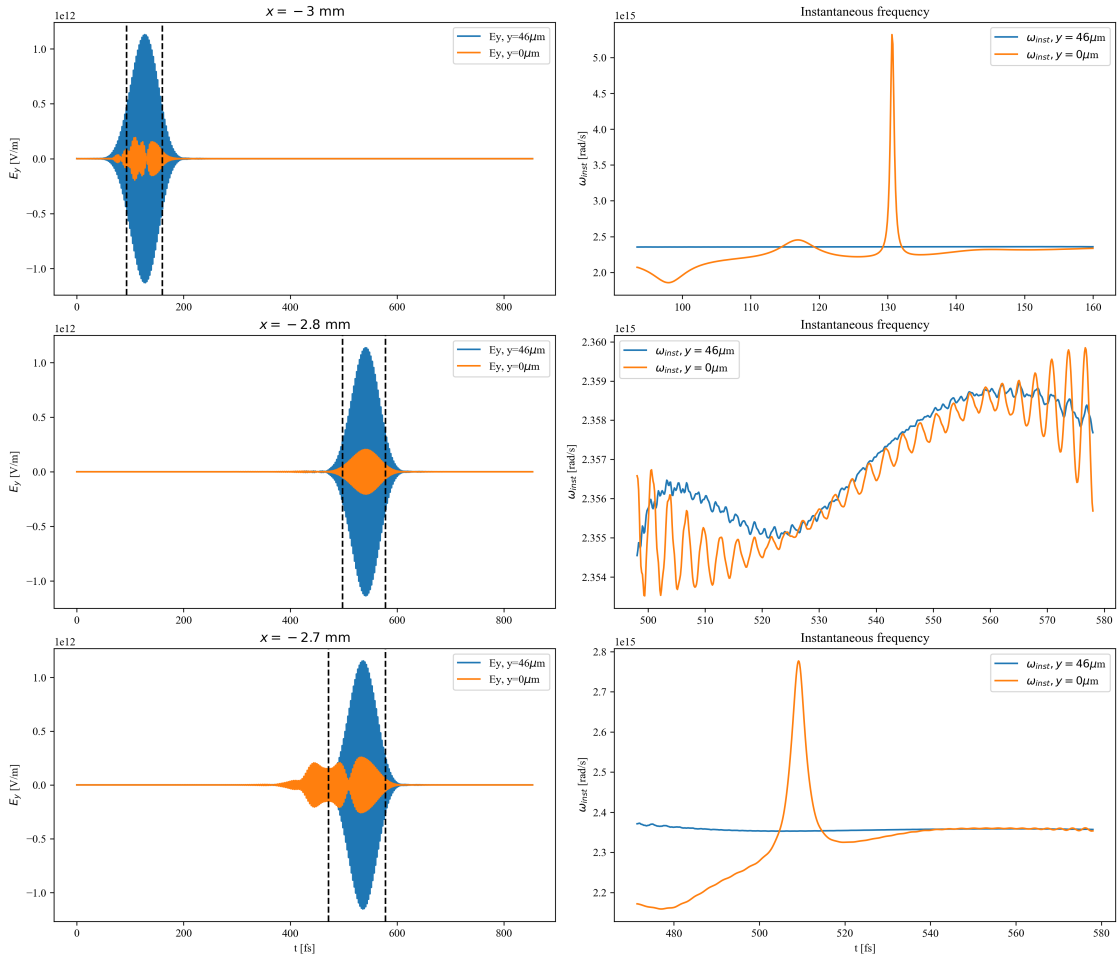

Figure 6.27: Left pannel: Transverse electric field at the position of the "hole" in orange in blue transverse electric field at the position of the maximum of one of the lobes. Right pannel corresponding instantaneous frequency. The differnet row are the different position represented by the red dashed line in FIG .6.12.

Appendix D: Self interference spatio spectral phase modulation

As discussed above, the large phase modulation cannot be explained by a Kerr-like mechanism, since the observed chirp is negative and the high-intensity regions do not exhibit significant phase modulation due to the low local intensity and we recover the same longitudinal field in vacuum simulation. The explanation lies elsewhere: this effect arises from spatio-spectral coupling, which is intrinsically linked to the frequency-dependent nature of light propagation.

Before analyzing the specific mechanisms behind this modulation, we first recall some fundamental facts about the nature of light propagation, which are often overlooked when discussing femtosecond pulses.

Consider a Gaussian laser beam. Its expression at a distance z from the focus in the (r, ω) domain is:

$$E(r,z,\omega) = e^{-\frac{(\omega-\omega_0)^2}{\Delta\omega^2}} e^{i\omega\frac{z}{c}} \int_{-\infty}^{+\infty} \exp\left[-\frac{k_\perp^2 w_0^2}{4} - i\left(\frac{k_\perp^2 zc}{2\omega} - k_\perp r\right)\right] dk_\perp. \tag{6.15}$$

Here, we used the paraxial approximation for the free-space propagation kernel:

$$k_z z \simeq \frac{\omega z}{c} - \frac{k_\perp^2 zc}{2\omega}. \tag{6.16}$$

The explicit evaluation of the integral gives:

$$E(r,z,\omega) = \sqrt{\pi}\, \frac{\exp\left(-\frac{r^2\omega}{w_0^2\omega+2icz}\right)}{\sqrt{\frac{w_0^2\omega+2icz}{\omega}}} e^{-\frac{(\omega-\omega_0)^2}{\Delta\omega^2}} e^{i\omega\frac{z}{c}}. \tag{6.17}$$

For our purposes, we can neglect the denominator, since in the regime $w_0^2\omega \gg cz$ this term only contributes a Gouy-like phase, which is approximately constant for our parameters. The expression then simplifies to:

$$E(r,z,\omega) = \exp\left[-\frac{r^2\omega^2 w_0^2}{w_0^4\omega^2 + 4z^2c^2} + i\frac{2r^2\omega zc}{w_0^4\omega^2 + 4z^2c^2}\right] e^{-\frac{(\omega-\omega_0)^2}{\Delta\omega^2}} e^{i\omega\frac{z}{c}}. \tag{6.18}$$

Now, replacing ω by ω_0 and noting that

$$\frac{z^2}{z_r^2} = \frac{4z^2c^2}{w_0^4\omega_0^2}, \tag{6.19}$$

We recover the standard Gaussian beam expression:

$$E(r,z,\omega) = \exp\left[-\frac{r^2}{w(z)^2} + ik_0\frac{r^2}{2R(z)}\right] e^{-\frac{(\omega-\omega_0)^2}{\Delta\omega^2}} e^{i\omega\frac{z}{c}}. \tag{6.20}$$

This simple example illustrates that the Gaussian beam expression is merely an approximation, which is typically valid because the signal is narrowband in frequency; under such conditions, the approximation usually has negligible consequences. Furthermore, it demonstrates that even for the simplest Gaussian beam, spatio-spectral coupling is intrinsic to polychromatic light propagation. In general, no light pulse can be strictly written as [14]

$$E(r,\omega) = f(r)\, g(\omega), \tag{6.21}$$

since group delay and frequency-dependent propagation effects are always present. In the following, we will show that in the regime of strongly aberrated beams, this separability assumption breaks down, providing an explanation for the phenomena observed in the simulations.

From free-space propagation to interference

To understand the phenomena observed in the simulations, we chose a simplified expression for the electric field that captures the essential physics of our more complex simulated beam. Since we observe a central low-energy region (the "hole") flanked by higher-energy side regions (the "lobes"), we define the following function, with its squared modulus shown in Fig. 6.28:

$$E(r) = \left(1 - \alpha\, \mathrm{e}^{\mathrm{i}\beta r}\right) \mathrm{e}^{-\frac{r^2}{w_0^2}} \tag{6.22}$$

This function is chosen not only because it preserves the observed symmetry, a central "hole" with surrounding "lobes",but also because the added phase term will be important in the analysis. Moreover, the simplicity of the expression allows for analytical tractability through the subsequent integral and derivatives.

We propagate the beam to a distance z by first taking the Fourier transform in $k_\perp$:

$$E(k_\perp) = \int_{-\infty}^{+\infty} \left(1 - \alpha\, \mathrm{e}^{\mathrm{i}\beta r}\right) \mathrm{e}^{-\mathrm{i}k_\perp r - \frac{r^2}{w_0^2}}\, dr \tag{6.23}$$

After removing the constant amplitude factor $\sqrt{\pi}w_0$, this yields:

$$E(k_\perp) = \mathrm{e}^{-\frac{w_0^2 k_\perp^2}{4}} - \alpha\, \mathrm{e}^{-\frac{w_0^2 k_\perp^2}{4} + \frac{w_0^2 \beta k_\perp}{2} - \frac{w_0^2 \beta^2}{4}} \tag{6.24}$$

The propagated field can then be written as (ignoring the $i\omega t$ term):

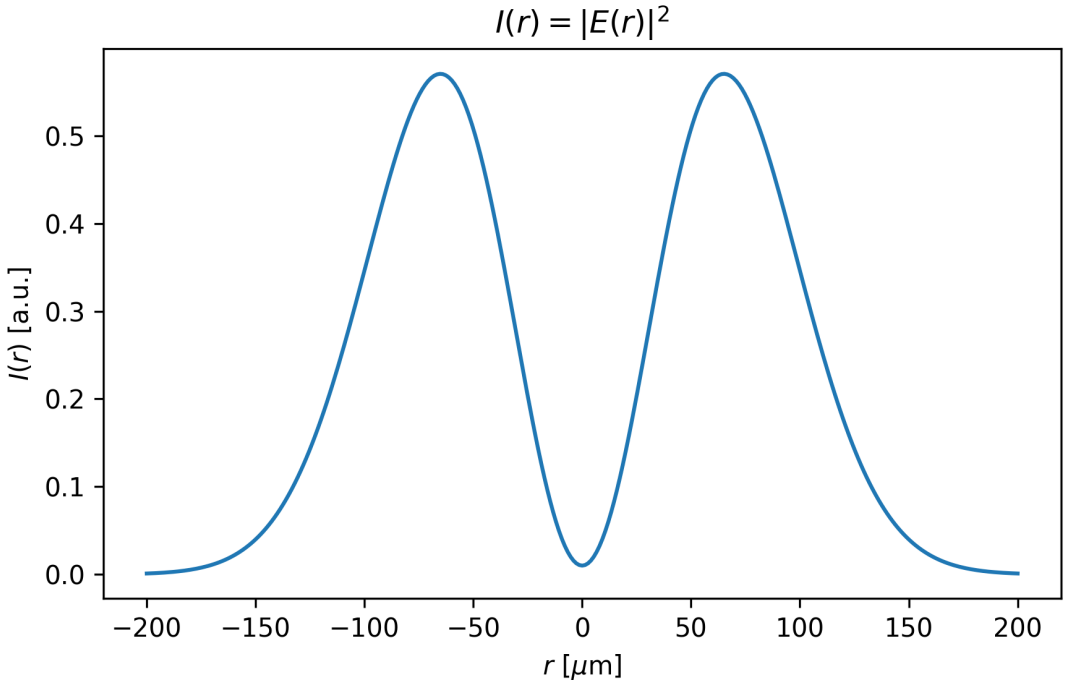

Figure 6.28: Spatial intensity of equation (6.22).

$$E(r,z,\omega) = \int_{-\infty}^{+\infty} \left(\mathrm{e}^{-\frac{w_0^2 k_\perp^2}{4}} - \alpha\, \mathrm{e}^{-\frac{w_0^2 k_\perp^2}{4} + \frac{w_0^2 \beta k_\perp}{2} - \frac{w_0^2 \beta^2}{4}} \right) \mathrm{e}^{\mathrm{i} r k_\perp - \frac{\mathrm{i} c z k_\perp^2}{2\omega}}\, dk_\perp \tag{6.25}$$

Which simplifies to (removing Gouy like phase term):

$$E(r,z,\omega) = \mathrm{e}^{-\frac{r^2 \omega}{w_0^2 \omega + 2icz}} - \alpha\, \mathrm{e}^{\frac{(2ir - w_0^2\beta)^2 \omega}{4(w_0^2\omega + 2icz)} - \frac{w_0^2\beta^2}{4}} \tag{6.26}$$

We now focus on the phase of this field to understand how the group delay varies with frequency, which leads to replicas of the pulse interfering with themselves and generating SISSPM.

Separating the two contributions on the right-hand side (RHS), we write:

$$E(r,z,\omega) = E_1(r,z,\omega) - \alpha\, E_2(r,z,\omega),$$

and focus on the spectral phase:

$$\phi(r,\omega) = \arg\left[E(r,z,\omega)\right].$$

We define the common denominator:

$$\Delta(\omega) \equiv w_0^4 \omega^2 + 4c^2 z^2.$$

First contribution: Gaussian background

Amplitude:

$$A(r,\omega) = \exp\left(-\frac{r^2 w_0^2 \omega^2}{\Delta} \right)$$

Phase:

$$\Phi_1(r,\omega) = \frac{2cz\omega r^2}{\Delta}$$

Second contribution: modulated replica

Amplitude:

$$B(r,\omega) = \exp\left[-\frac{w_0^2\beta^2}{4} + \frac{w_0^2\omega^2}{\Delta}\left(-r^2 + \frac{w_0^4\beta^2}{4} \right) - \frac{2cz\omega r w_0^2 \beta}{\Delta} \right]$$

Phase:

$$\Phi_2(r,\omega) = \frac{2cz\omega}{\Delta}\left(r^2 - \frac{w_0^4\beta^2}{4} \right) - \frac{\omega^2 r w_0^4 \beta}{\Delta}$$

Total spectral phase

$$\phi(r,\omega) = \arg\left[E_1(r,z,\omega) - \alpha\, E_2(r,z,\omega)\right]$$

The total phase of the field is therefore

$$\phi(r,\omega) = \arctan\left(\frac{A(r,\omega)\,\sin\big(\Phi_1(r,\omega)\big) - \alpha\, B(r,\omega)\,\sin\big(\Phi_2(r,\omega)\big)}{A(r,\omega)\,\cos\big(\Phi_1(r,\omega)\big) - \alpha\, B(r,\omega)\,\cos\big(\Phi_2(r,\omega)\big)}\right)$$

Now we want to understand how the group delay varies with r and ω, which requires computing the derivative. Since the derivative is particularly cumbersome, we used an online derivative calculator, and the full expression is provided in the appendix. Here, we present the group delay defined as

$$\tau(r,\omega) = \frac{d\phi}{d\omega}.$$

In Fig. 6.29 we show two representative cases. In the first case (left panel), the parameters are $w_0 = 100 \times 10^{-6}$ m, $z = 10$ μm, $\beta = 2 \times 10^4$ rad/m, and $\alpha = 0.9$. In the second case (right panel), we have $w_0 = 100 \times 10^{-6}$ m, $z = 10$ mm, $\beta = 2 \times 10^4$ rad/m, and $\alpha = 0.999$.

The first observation is that as α and z increase, the maximum group delay also increases. Furthermore, from the figure we see that the largest values of the group delay coincide with the position of the "hole" as z increases.

In Fig. 6.30, we present the effective group delay dispersion in seconds, $GDD \cdot \Delta\omega$, for both cases. In the second case, the difference in arrival time between high and low frequencies reaches approximately 60 fs. As with the group delay, the largest values are located at the central "hole." This matches our Smilei vacuum simulation, where the chirp is negative.

To generalize, consider any complex beam of the form:

$$E(r,\omega) = |E(r,\omega)|e^{i\phi(r,\omega)}.$$

The group delay can then be expressed as:

$$\tau(r,\omega) = \frac{\mathrm{Im}\left[\left(\frac{\partial E}{\partial \omega} - \frac{\partial |E|}{\partial \omega}e^{i\phi}\right)E^*\right]}{|E(r,\omega)|^2}. \tag{6.27}$$

Separating the real and imaginary parts, we write

$$E(r,z,\omega) = [A\cos\phi_1 - \alpha B\cos\phi_2] + i[A\sin\phi_1 - \alpha B\sin\phi_2], \tag{6.28}$$

so that the squared modulus is

$$|E(r,z,\omega)|^2 = A^2 + \alpha^2 B^2 - 2\alpha AB\cos(\phi_1 - \phi_2). \tag{6.29}$$

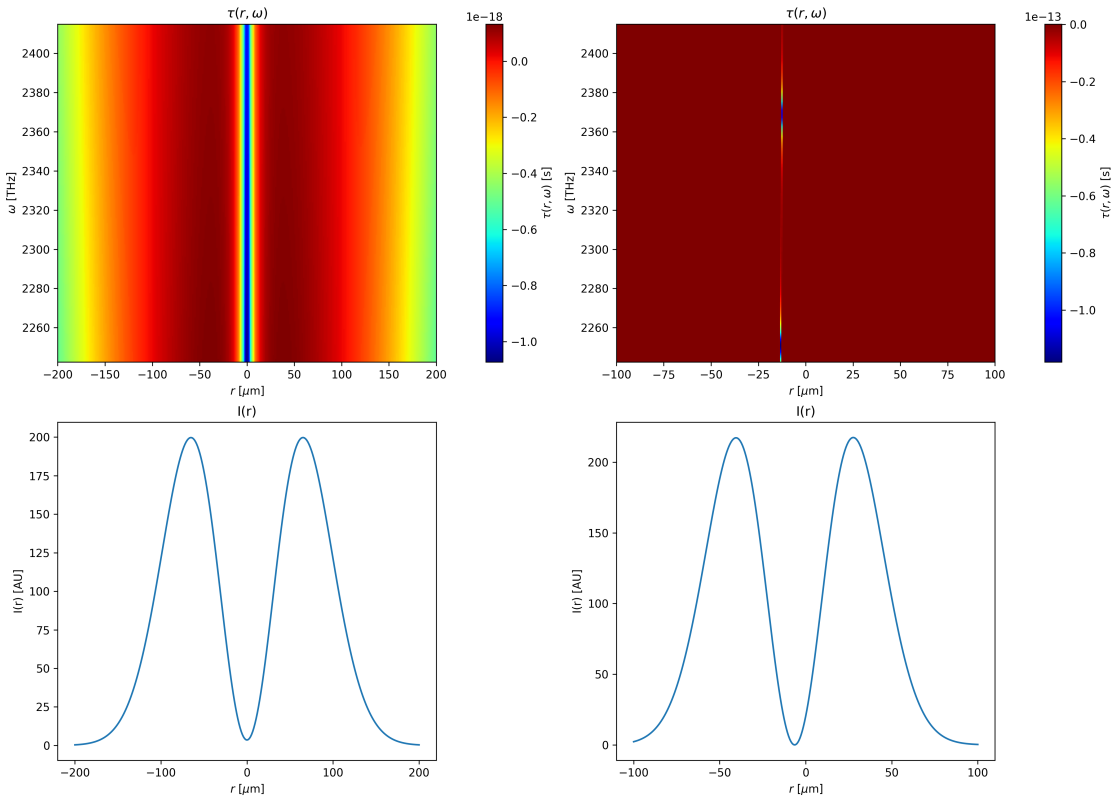

Figure 6.29: Top row: Group delay for the two cases. Bottom row: corresponding spatial intensity.

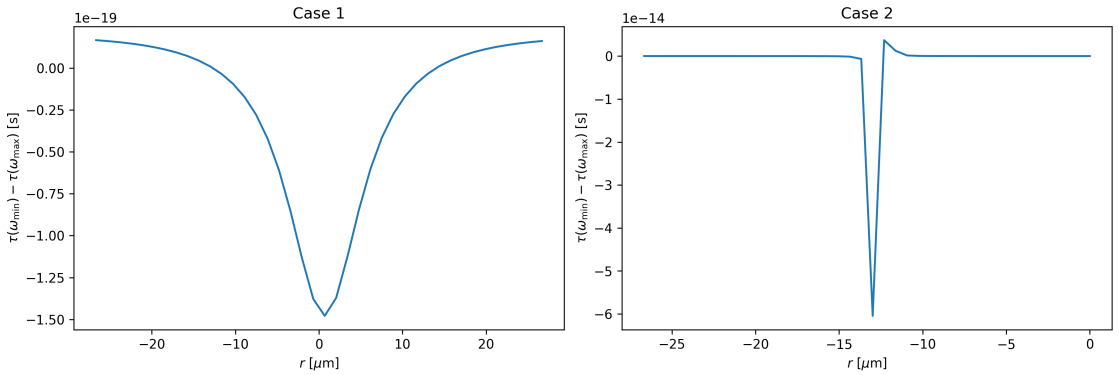

Figure 6.30: Effective group delay dispersion ($GDD \cdot \Delta\omega$) for the two cases.

We can already make an important observation: the squared modulus contains an interference term. If $A \sim B$, $\alpha \sim 1$, and $\phi_1 - \phi_2 \sim 0$, the denominator in equation (6.27) becomes small, leading to large frequency-dependent group delays.

To analyze the conditions for this to occur, consider the interference term $\phi_1 - \phi_2$ at $r = 0$ near the central "hole":

$$\phi_1 - \phi_2 = \frac{cz\beta^2\omega w_0^4}{2(\omega^2 w_0^4 + 4c^2 z^2)}.$$

Introducing the Rayleigh-length-dependent factor:

$$\frac{z^2}{z_r(\omega)^2} = \frac{4c^2 z^2}{w_0^4 \omega^2},$$

We can rewrite:

$$\phi_1 - \phi_2 = \frac{cz\beta^2}{2\omega\left(1 + \frac{z^2}{z_r(\omega)^2}\right)}.$$

In the regime $\omega \gg cz\beta^2$, we have $\cos(\phi_1 - \phi_2) \sim 1$ near the central "hole," satisfying the first condition.

Next, we evaluate when $A \sim B$. Near $r = 0$:

$$A = 1, \quad B = \exp\left[-\frac{w_0^2\beta^2}{4} + \frac{w_0^2\beta^2}{4(1 + z^2/z_r(\omega)^2)}\right],$$

so that

$$B \in \left[\exp\left(\frac{-w_0^2\beta^2 \frac{z^2}{z_r(\omega)^2}}{4(1 + z^2/z_r(\omega)^2)}\right), 1\right].$$

Thus, the condition $A \sim B$ is satisfied if

$$w_0^2\beta^2 \ll 1 \quad \text{and/or} \quad z \ll z_r.$$

In general, the effect occurs whenever $A \sim B$, $\alpha \sim 1$, and $\beta \neq 0$.

For $r \neq 0$, the phase difference becomes

$$\phi_1 - \phi_2 = \frac{cz\beta^2}{2\omega(1 + z^2/z_r(\omega)^2)} - \frac{r\beta}{1 + z^2/z_r(\omega)^2},$$

producing an oscillatory term. As a result, this effect is strongest near the central "hole," provided $w_0\beta < 1$ (i.e., the beam width is smaller than the oscillation period).

Finally, we can quantify when this effect becomes significant. For femtosecond-scale pulses, the condition is

$$\max(\tau(r,\omega)) > \text{fs}.$$

Assuming $z \ll z_r(\omega)$ and $\omega \gg cz\beta^2$, the group delay at $r = 0$ simplifies to:

$$\tau(\omega) = \frac{-1}{2}\frac{cz\beta^2}{(\alpha-1)\omega^2}.$$

Since this expression is always positive in magnitude, it systematically leads to a negative chirp,i.e.the highest frequencies arrive first. Thus, near the "hole":

$$\max(\tau(r,\omega)) = \left|\frac{1}{2}\frac{cz\beta^2}{(\alpha-1)\omega_{\text{min}}^2}\right| > \text{fs}. \tag{6.30}$$

Generalization and interpretation

One might argue that our previous example works because the function we chose was engineered to satisfy the conditions necessary for frequency-dependent group delays near the central "hole." However, we can make more general statements starting from a broader, more universal expression.

The key observation from our example is that the large frequency-dependent group delays near the central hole are created by destructive interference. Therefore, the general expression for the electric field must take the form:

$$E(r,\omega) = \sum_j E_j e^{i\phi_j}.$$

At a minimum, the field needs at least two main contributors:

$$E(r,\omega) = E_1 e^{i\phi_1} + E_2 e^{i\phi_2}. \tag{6.31}$$

Since the phenomenon arises from destructive interference, the first necessary condition is that at the "hole" location $r = r_{\text{hole}}$:

$$\left.\frac{\partial|E(r,\omega)|^2}{\partial r}\right|_{r=r_{\text{hole}}} = 0 \quad \text{and} \quad \left.\frac{\partial^2|E(r,\omega)|^2}{\partial r^2}\right|_{r=r_{\text{hole}}} > 0.$$

That is, the electric field must be a minimum at the hole. Using expression (6.31), the squared modulus reads:

$$|E|^2 = (E_1^2 + E_2^2)\left(1 + \frac{2E_1E_2}{E_1^2+E_2^2}\cos(\phi_1-\phi_2)\right).$$

For the minimum condition to be satisfied, it requires:

$$\left|\frac{2E_1E_2}{E_1^2+E_2^2}\right| = |\alpha| \sim 1 \quad \text{and} \quad \lim_{r\to r_{\text{hole}}}(\phi_1-\phi_2) \ll 1 \sim 0 \quad \text{or} \quad \lim_{r\to r_{\text{hole}}}(\phi_1-\phi_2) \sim \pi,$$

where the choice between 0 or π depends on the sign of the contrast α; in our case, $\alpha < 0$.

A general expression that satisfies the phase condition near $r = r_{\text{hole}}$ can be written as a polynomial expansion:

$$\phi_1 - \phi_2 = f_0(\omega) + f_1(\omega)\beta(r - r_{\text{hole}}) + f_2(\omega)\gamma(r - r_{\text{hole}})^2 + \dots \tag{6.32}$$

Here, $f_i(\omega)$ are functions of frequency, $f_0(\omega) \ll 1$, and β, γ are coefficients of the spatial aberrations.

Comment: For SISSPM to occur, the initial laser beam must possess spatial aberrations encoded in β and γ. Without them, the minimum condition at the hole is violated.

The second set of requirements comes from the numerator of equation (6.27). To observe SISSPM, the group delay must be frequency-dependent:

$$\left|\text{Im}\left[\left(\frac{\partial E(\omega)}{\partial\omega} - \frac{\partial|E(\omega)|}{\partial\omega}e^{i\phi(\omega)}\right)E^*(\omega)\right]\right| = f(\omega) > 0. \tag{6.33}$$

Assuming the minimum condition is satisfied, i.e., at the hole:

$$|E| = \sqrt{E_1^2 + E_2^2 + 2E_1E_2\cos(\phi_1-\phi_2)} \ll \sqrt{E_1^2+E_2^2}, \quad and \quad \frac{\partial|E|}{\partial\omega} \sim 0$$

At the hole the modulus depends weakly on frequency which is reasonable since the field is minimal at the "hole" and we don't consider the spectrum. The expression simplifies to:

$$\left|\frac{\partial\phi_1}{\partial\omega}E_1^2 + \frac{\partial\phi_2}{\partial\omega}E_2^2 + E_1E_2\left(\frac{\partial\phi_1}{\partial\omega}+\frac{\partial\phi_2}{\partial\omega}\right)\cos(\phi_1-\phi_2)\right| = f(\omega) > 0. \tag{6.34}$$

If we consider

$$\cos(\phi_1-\phi_2) \simeq -\frac{E_1^2+E_2^2}{2E_1E_2},$$

the numerator reduces to:

$$\left|\frac{1}{2}\frac{\partial f_0(\omega)}{\partial\omega}(E_1^2-E_2^2)\right| = f(\omega) > 0. \tag{6.35}$$

Hence, SISSPM is guaranteed as long as $|\partial_\omega f_0| = f(\omega) > 0$ and $|E_1| \neq |E_2|$, i.e., the destructive interference does not completely cancel the field at the hole.

To extract a strength parameter for SISSPM, consider the minimal initial field:

$$E(r) = E_1 + E_2 e^{i\beta r}.$$

Focusing on the second term with $E_2 = 1$, the Fourier transform in k-space gives:

$$E_2(k_\perp) = \mathscr{F}[e^{i\beta r}] = \delta(k_\perp - \beta).$$

Under the paraxial approximation, the propagation integral becomes:

$$E_2(r, z, \omega) = e^{i\omega z/c} \int_{-\infty}^{+\infty} \delta(k_\perp - \beta) e^{irk_\perp - iczk_\perp^2/(2\omega)} dk_\perp, \tag{6.36}$$

yielding

$$E_2(r, z, \omega) = \exp\left(i\beta r - i\frac{cz\beta^2}{2\omega}\right). \tag{6.37}$$

Thus, the frequency-dependent term is

$$f_0(\omega) = \frac{cz\beta^2}{2\omega}. \tag{6.38}$$

Finally, if we include a quadratic spatial aberration in the second term:

$$E_2(r) = e^{i\gamma r^2},$$

after paraxial propagation we get (neglecting Gouy like phase and amplitude terms):

$$E_2(r, z, \omega) = \exp\left[\frac{ir^2\gamma\omega}{\omega + 2zc\gamma}\right], \tag{6.39}$$

so that the phase difference can be expressed as

$$\phi_1 - \phi_2 = f_1(\omega)\gamma(r - r_{\text{hole}})^2,$$

In other words, $f_0(\omega) = 0$ and the necessary condition $|\partial_\omega f_0| = f(\omega) > 0$ is violated. In general, this argument also holds when considering higher-order polynomial terms in r, such as $r^3, r^4, \ldots$. Thus, the last necessary condition to observe SISSPM is that the phase difference $\phi_1 - \phi_2$ possesses an effective spatial phase tilt β.

If we summarize the previous conditions, SISSPM will occur only if at least two comparable contributions interfere to produce a near[2] but non-zero minimum via destructive interference at the position of the "hole." The second condition is that the laser beam must also possess an effective spatial phase tilt. Lastly, the sign of the chirp

[2]Here we do not mean that the field amplitude at the "hole" position vanishes. In principle, it can be arbitrarily large; what matters is that it is much smaller than the neighboring maxima, i.e., that the interference contrast parameter satisfies $|\alpha| \sim 1$.

is determined by the sign of $E_1 - E_2$ in the numerator. Intuitively, this effect can be understood as follows. In vacuum, the dispersion relation allows different frequency components to propagate at slightly different angles: higher frequencies correspond to larger allowed propagation cones. If one considers a beam with an initial spatial phase tilt propagating over a distance Z, the higher-frequency components naturally arrive earlier. Additionally, destructive interference reinforces this temporal ordering. As a result, the spatial phase tilt couples to the frequency components through the propagation kernel $k_z z$, reflecting the frequency-dependent nature of free-space propagation.

Some strength parameters for SISSPM can therefore be extracted from the previous considerations.

$$\Lambda = \frac{w_0|\beta|}{1-|\alpha|} \tag{6.40}$$

This first parameter indicates whether the effect can occur at a given longitudinal position z. The effect becomes more likely as Λ increases.

$$\Gamma = \frac{c(z-z_0)\beta^2}{(1-|\alpha|)\Delta\omega} \tag{6.41}$$

This second parameter is related to propagation, with z_0 being the first longitudinal plane where all SISSPM conditions are simultaneously satisfied. The longer these conditions remain satisfied, the larger Γ becomes, since $\Gamma \propto z$. The last term to comment on is $\Delta\omega$: the larger the initial frequency bandwidth, the greater the potential group-delay dispersion $GDD(\omega)$ that can develop.

This also explains why the phenomenon does not occur for beams such as ideal Laguerre–Gaussian or Hermite–Gaussian modes, which possess symmetry around a central intensity minimum. These modes lack any effective spatial phase tilt, and the electric field cancels perfectly at the center of the "hole." As a result, the required frequency-dependent phase offset $f_0(\omega)$ and the non-zero electric field conditions are absent.

In Fig. 6.31, we present several quantities extracted from the vacuum simulation in order to test our model. Since we do not have direct access to the value of α,as this would require assuming a specific functional form, we instead define an effective parameter:

$$\alpha_{\text{eff}} = \frac{I_{\max} - I_{\text{hole}}}{I_{\max} + I_{\text{hole}}}.$$

This choice is reasonable, as it avoids selecting a specific functional form and is expected to remain proportional to α at all times.

The top-left panel of Fig. 6.31 shows α_{eff} as a function of the longitudinal position. We observe that it decreases significantly with propagation, which is expected since the laser is focusing. The top-right panel shows β, defined as:

$$\beta = \sqrt{\beta_x^2 + \beta_y^2}.$$

Our model is two-dimensional, whereas the simulation is three-dimensional, which motivates this definition. The parameter β is extracted from a Zernike polynomial regression. Since β varies significantly due to imperfections in the regression, and since the essential requirement for SISSPM is that $\beta \neq 0$,ensuring $|\partial_\omega f_0| > 0$ at all times,we choose to fix $\beta = 1$ when computing $\Gamma(t)$ and $\Lambda(t)$. Furthermore, the dominant mechanism leading to SISSPM is destructive interference.

The bottom-left panel represents

$$\omega(t = t_{\mathrm{last}}) - \omega(t = t_{\mathrm{first}}),$$

that is, the instantaneous frequency at the end of the analysis window minus the instantaneous frequency at the beginning of the window. This quantity is intentionally not defined as a max–min difference, since the true maxima and minima of the instantaneous frequency are typically dominated by numerical and physical singularities when the intensity becomes very small, as shown in Fig. 6.32.

This choice is physically motivated: Eq. (6.27) shows that when the intensity approaches zero, the phase becomes ill-defined, and the instantaneous frequency loses physical meaning. The window is therefore defined only over the region where the field amplitude exceeds 10% of its maximum value, as indicated by the black dashed line in Fig. 6.32.

In the bottom-right panel of Fig. 6.31, we observe that the theoretical prediction closely matches the simulation results: $\Gamma(t)$ reproduces the same longitudinal evolution as the instantaneous frequency differential.

Finally, the decrease of the instantaneous frequency differential around $-2.2\ \mathrm{mm}$ is a direct consequence of this windowing procedure. Because the analysis window is defined by a fixed relative amplitude threshold, the effective temporal window shrinks as the pulse structure evolves. If the full temporal domain were considered, the instantaneous frequency differential would very likely remain approximately constant beyond $-2.2\ \mathrm{mm}$. However, since the phase is not physically meaningful at low intensities, the use of such a window is necessary.

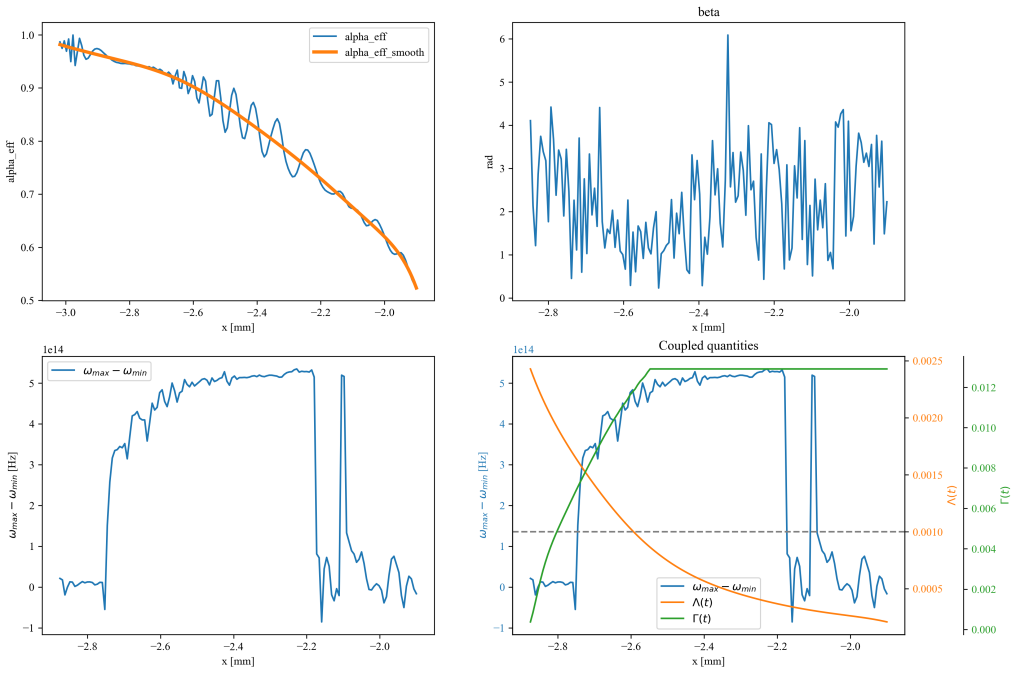

Figure 6.31: Top-left panel: α_{eff} as a function of the longitudinal position. Top-right panel: β as a function of the longitudinal position. Bottom-left panel: instantaneous frequency differential defined as $\omega(t_{\text{last}}) - \omega(t_{\text{first}})$ within the analysis window. Bottom-right panel: $\Gamma(t)$, $\Lambda(t)$, and the instantaneous frequency differential as functions of the longitudinal position. The dashed line indicates the position at which the SISSPM condition is no longer satisfied, corresponding to $\alpha_{\text{eff}} = 0.91$.

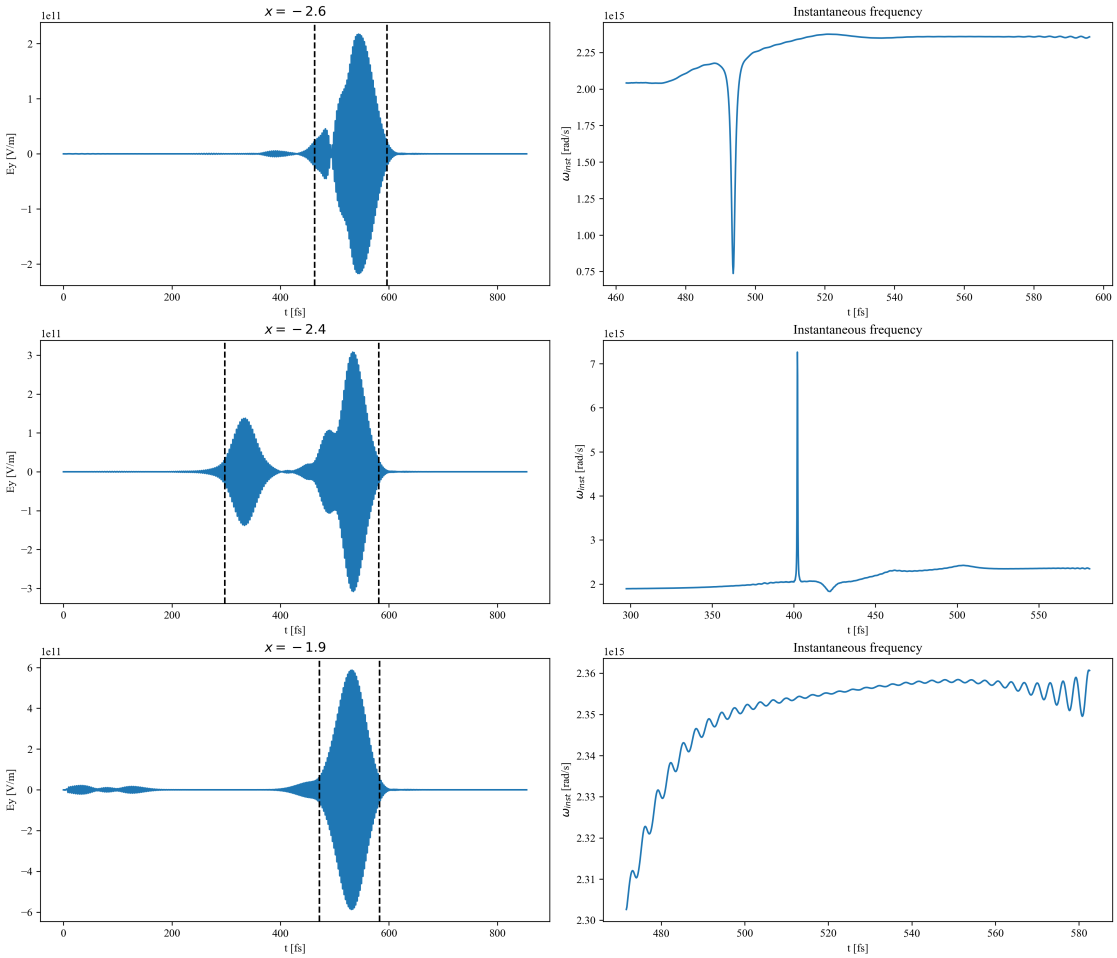

Figure 6.32: Left panel: transverse electric field at the position of the “hole” in the vacuum simulation. Right panel: corresponding instantaneous frequency. Each row corresponds to a different longitudinal position. The black dashed line indicates the analysis window defined by the 10% intensity threshold.

We can now attempt to explain the origin of the large longitudinal electric fields observed in the simulation. These fields are most likely a *secondary consequence of SISSPM*. The strong spatio-spectral phase modulation causes the pulse to extend well beyond its nominal temporal duration τ_{FWHM} (see Fig. 6.26). As a result, a non-negligible fraction of the transverse field E_y persists outside the main pulse envelope and reaches the spatial location of the central intensity minimum (the "hole").

At this position, the effective transverse size of the beam becomes extremely small. In such a regime, significant longitudinal electric-field components can arise. As shown by Lax *et al.* [15] and later by Chen *et al.* [16], when the term $\nabla(\nabla \cdot \mathbf{E})$ is retained in the wave equation, strong longitudinal electric fields naturally emerge. The magnitude of this longitudinal component scales as

$$f \sim (w_0 k)^{-1},$$

so that a strong enhancement is expected when the beam waist w_0 becomes very small. Since w_0 is extremely small beyond the hole location, this mechanism provides a plausible explanation for the large longitudinal electric fields observed in the simulation. This hypothesis would require dedicated numerical verification for the specific parameters considered here.

Based on these considerations, the first two spikes in the longitudinal electric field are likely numerical artifacts arising from boundary condition effects in Smilei, whereas the last spike is likely physical. Its precise magnitude, however, remains uncertain and warrants further investigation. A simple argument to verify that the values of E_x obtained from the Smilei simulation are physically consistent is presented in Fig. 6.33. Using equations (2) and (3) in Ref. [16], we computed the longitudinal and transverse electric fields at the position $x = -2.7\ \mathrm{mm}$, starting from the initial laser field shown in Fig. 6.6.

Note that the Smilei electric field is obtained from a probe located at $x = 155.3\ \mu\mathrm{m}$, as shown in Fig. 6.26. As seen in Fig. 6.33, the electric fields from the Smilei simulation closely match the fields obtained by propagation using equations (2) and (3) in [16]. This provides a simple validation that the Smilei E_x values are reliable.

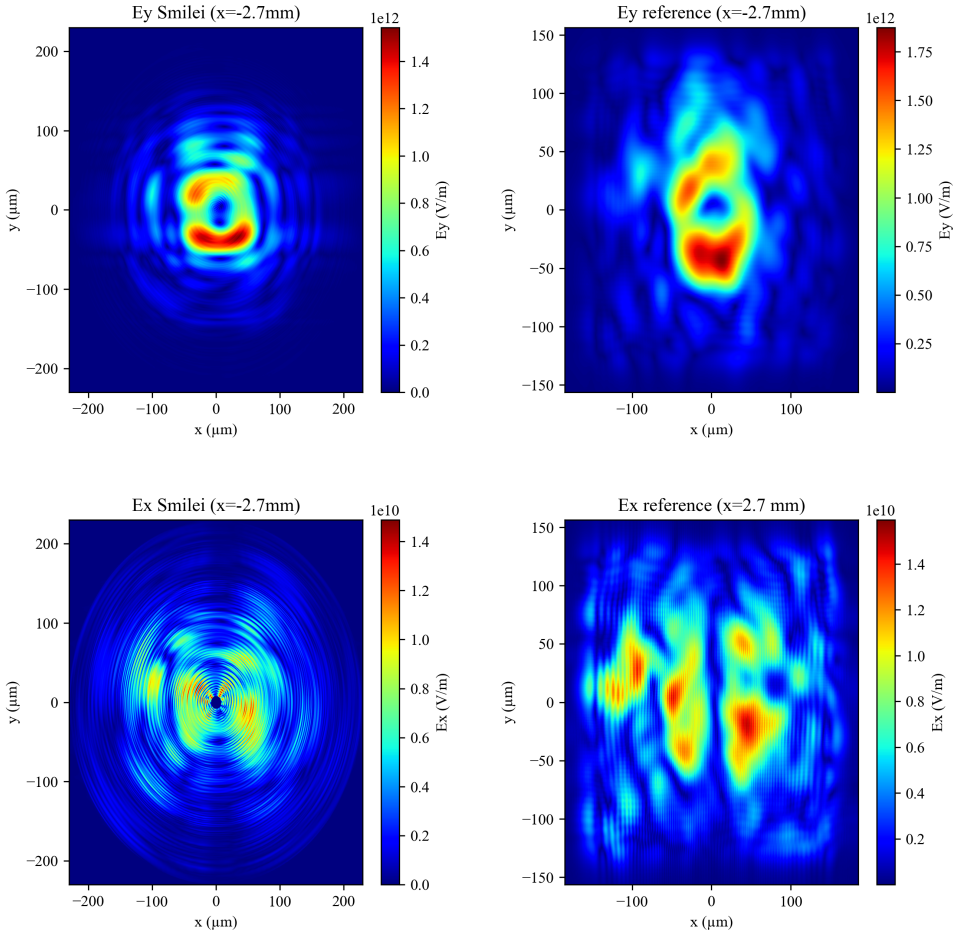

Figure 6.33: Top left: transverse electric field (E_y) at $x = -2.7$ mm from the Smilei probe located at $x = 155.3$ μm in Fig. 6.26. Top right: transverse electric field (E_y) computed at $x = -2.7$ mm using the propagation of the field in Fig. 6.6. Bottom left: longitudinal electric field (E_x) at $x = -2.7$ mm from the Smilei probe. Bottom right: longitudinal electric field (E_x) obtained by propagation using the field in Fig. 6.6.

Appendix E: Full expression of the group delay derivative

For completeness, we provide the full derivative of the spectral phase $\tau(r,\omega) = \frac{d\phi}{d\omega}$ in a structured form.

$$D = w_0^4\omega^2 + 4c^2z^2 \tag{6.42}$$

$$A = 2rw_0^4\beta\omega + cw_0^4z\beta^2 - 4cr^2z \tag{6.43}$$

$$G = \exp\left(\frac{cw_0^2z\beta(2r\omega + cz\beta)}{D}\right) \tag{6.44}$$

$$\Phi_1 = \frac{\omega A}{2D}, \qquad \Phi_2 = \frac{2cr^2z\omega}{D} \tag{6.45}$$

$$U = \frac{A}{2D} + \frac{rw_0^4\beta\omega}{D} - \frac{w_0^4\omega^2 A}{D^2} \tag{6.46}$$

$$V = \frac{2crw_0^2z\beta}{D} - \frac{2cw_0^6z\beta\omega(2r\omega + cz\beta)}{D^2} \tag{6.47}$$

$$W = \frac{2cr^2z}{D} - \frac{4cr^2w_0^4z\omega^2}{D^2} \tag{6.48}$$

$$S = \alpha\sin\Phi_1 + G\sin\Phi_2 \tag{6.49}$$

$$C = \alpha\cos\Phi_1 - G\cos\Phi_2 \tag{6.50}$$

$$T_1 = \frac{\alpha U\cos\Phi_1 + VG\sin\Phi_2 + WG\cos\Phi_2}{C} \tag{6.51}$$

$$T_2 = \frac{S\left(-\alpha U\sin\Phi_1 + WG\sin\Phi_2 - VG\cos\Phi_2\right)}{C^2} \tag{6.52}$$

$$\frac{d\phi}{d\omega} = -\frac{T_1 - T_2}{\frac{S^2}{C^2} + 1} \tag{6.53}$$

References

[1] SMILEI Developers. *Azimuthal modes decomposition*. https://smileipic.github.io/Smilei/Understand/azimuthal_modes_decomposition.html. Accessed: 2026-03-02. 2024.

[2] Imene Zemzemi. "High-performance computing and numerical simulation for laser wakefield acceleration with realistic laser profiles". NNT: 2020IPPAX111. PhD thesis. Institut Polytechnique de Paris, 2020.

[3] W. Lu et al. "Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime". In: *Physical Review Accelerators and Beams* 10 (2007), p. 061301.

[4] X. L. Xu et al. "Phase-Space Dynamics of Ionization Injection in Plasma-Based Accelerators". In: *Physical Review Letters* 112.3 (201), p. 035003. doi: 10.1103/PhysRevLett.112.035003.

[5] Maxence Thévenet et al. *LASY: LAser Manipulations Made eaSY*. Mar. 2024. doi: 10.48550/arXiv.2403.12191. arXiv: 2403.12191 [physics]. (Visited on 03/05/2026).

[6] O. E. Martinez. "Pulse distortions in tilted pulse schemes for ultrashort pulses". In: *Optics Communications* 59.3 (1986), pp. 229–232.

[7] C. Dorrer, E. M. Kosik, and I. A. Walmsley. "Spatio-temporal characterization of ultrashort optical pulses using two-dimensional shearing interferometry". In: *Applied Physics B: Lasers and Optics* 74.suppl. (2002), pp. 209–219.

[8] Z. Bor et al. "Femtosecond pulse front tilt caused by angular dispersion". In: *Optical Engineering* 32.10 (1993), pp. 2501–2503.

[9] J. Hebling. "Derivation of pulse-front tilt caused by angular dispersion". In: *Optical and Quantum Electronics* 28 (1996), pp. 1759–1763.

[10] Selcuk Akturk et al. "The General Theory of First-Order Spatio-Temporal Distortions of Gaussian Pulses and Beams". In: *Optics Express* 13.21 (2005), pp. 8642–8661. doi: 10.1364/OPEX.13.008642.

[11] Erik Löfquist et al. "Correlated x-ray and electron beam steering by pulse-front tilt in laser wakefield acceleration". In: *Physical Review Accelerators and Beams* 28.10 (2025), p. 102802.

[12] Maxence Thévenet et al. "Pulse front tilt steering in laser plasma accelerators". In: *Physical Review Accelerators and Beams* 22.07 (2019), p. 071301. doi: 10.1103/PhysRevAccelBeams.22.071301.

[13] Daniel E. Mittelberger et al. “Laser and electron deflection from transverse asymmetries in laser-plasma accelerators”. In: *Physical Review E* 100 (2019), p. 063208. doi: 10.1103/PhysRevE.100.063208.

[14] Miguel A. Porras. “Ultrashort pulsed Gaussian light beams”. In: *Physical Review E* 58.1 (1998), pp. 1086–1095. doi: 10.1103/PhysRevE.58.1086.

[15] Melvin Lax, William H. Louisell, and William B. McKnight. “From Maxwell to paraxial wave optics”. In: *Physical Review A* 11.4 (1975), p. 1365.

[16] Carl G. Chen et al. “Analyses of vector Gaussian beam propagation and the validity of paraxial and spherical approximations”. In: *Journal of the Optical Society of America A* 19.2 (2002), pp. 404–412.

Conclusion

Throughout this PhD, I focused on the development of advanced laser diagnostics and machine learning tools, with the ultimate objective of enabling future implementation of feedback loops on laser and plasma parameters to improve the quality and stability of electron beams in laser-plasma accelerators.

I participated in the PALLAS commissioning phase and took part in the campaign that led to the production of the first electrons on the platform. This campaign demonstrated that the PALLAS platform is capable of delivering electron beams with relatively good quality, despite the fact that some building blocks had not yet been implemented, such as beam pointing stabilization and quadrupoles. Further efforts are currently underway to achieve a tunable laser–plasma injector source capable of delivering high-quality electron beams.

I developed surrogate models (SM) that demonstrated excellent regression performance on low-fidelity simulations, offering the major advantage of extremely fast inference compared to full PIC simulations. Among the ML methods explored, multi-layer perceptrons (MLP) and Gaussian processes (GP) achieved the highest performance. I highlighted the importance of the input data distribution, showing that a uniform sampling of input parameters is essential for building efficient SM. A methodology was presented to use these models for predicting electron beam parameters, which could be extended to experimental data or more realistic simulations, ultimately supporting feedback loops on laser and plasma parameters for enhanced electron beam stability. Additionally, I extended the SM approach to predict the electron spectrum. While this model showed lower performance, it could be improved through increased model complexity.

In a second line of research, I employed a multispectral camera (MSCCD) to optimize laser characteristics. The MSCCD demonstrated strong performance in aligning the pulse compressor and diagnosing spatial chirp. Using both simulated and experimental data, I showed that the MSCCD is a reliable tool for retrieving the spatio-spectral phase via the Gerchberg–Saxton algorithm (GSA). Furthermore, I developed a machine learning–based algorithm that achieves comparable performance to GSA, with the advantages of faster computation and direct retrieval of the unwrapped phase. This approach could enable near–real-time spatio-spectral phase reconstruction in high-repetition-rate operation. I also demonstrated that, in combination with a spectral phase measurement device, the MSCCD could fully resolve the 3D electric field of the laser, except for its polarization state. Such techniques could eventually be coupled with spatio-spectral phase control mechanisms and implemented in a feedback loop to tailor and stabilize the electron beam quality produced by laser-plasma acceleration.

Thirdly, I developed a thermo-mechanical pipeline, extending the work of Leroux et al., capable of predicting the spatio-spectral amplitude and phase aberrations induced

by compressor heating. The code produced predictions in good agreement with experimental measurements. This study emphasized that, for laser-plasma injectors operating above 10 Hz, the choice of pulse compressor material is critical; otherwise, thermally induced aberrations prevent the generation of stable, high-quality electron beams.

Finally, I carried out realistic simulations using the methods developed above. First, these simulations validated the relevance of the ML approach, showing that even a laser beam with moderate aberrations can produce electron beams with parameters close to those predicted by ideal low-fidelity simulations. Second, I identified the physical mechanism responsible for the emergence of pulse-front tilt (PFT) inside a plasma, even when the laser beam is initially free of spatio-temporal aberrations. Third, I discovered a pulse-shaping effect called self-interference spatio-spectral phase modulation (SISSPM) and proposed a preliminary model to explain it.

Overall, the methods developed during this thesis provide a foundation for future researchers to implement feedback control systems based on machine learning and advanced laser diagnostics. Together with recent developments from other research teams, this work suggests that machine learning and computational tools will be essential to make practical laser-plasma accelerators a reality.