Université du Québec
Centre Énergie Matériaux Télécommunications
INRS Institut national de la recherche scientifique

MACHINE LEARNING OPTIMIZATION FOR IMPROVED LASER-DRIVEN ION BEAMLINES

OPTIMISATION PAR APPRENTISSAGE AUTOMATIQUE DES LIGNES DE FAISCEAUX D'IONS ACCÉLÉRÉS PAR INTERACTION LASER-MATIÈRE

Par
Elias Catrix

Thèse présentée pour l'obtention du grade de
Philosophiae Doctor, Ph.D.
en Sciences de l'énergie et des matériaux

Jury d'évaluation

Président du jury et Examinateur interne	François Vidal INRS-ÉMT, Varennes, Canada
Examinatrice externe	Félicie Albert Lawrence Livermore National Laboratory, Livermore, USA
Examinatrice externe	Charlotte Palmer Queen's University, Belfast, Northern Ireland, United Kingdom
Directeur de recherche	Patrizio Antici INRS-ÉMT, Varennes, Canada

À toutes celles et ceux qui rêvent et qui osent, pour que chacun puisse poursuivre ses ambitions, quel que soit son sexe, son origine ou sa nationalité. Que la science et la connaissance soient un terrain d'égalité et d'émancipation.

REMERCIEMENTS / ACKNOWLEDGEMENTS

Comment ne pas commencer par la personne qui partage ma vie sans jamais douter de la personne que je suis. Ninou, merci pour ton soutien indéfectible et ta bonne humeur communicatrice. De m'avoir inculqué la compassion, la patience et le partage qui te caractérisent tant. Merci infiniment de me laisser regarder Lyon, Arsenal, le Tennis, Le Heat de Miami, Sans Filet, Trashtalk, le WFC ou Basket Session même si je les regarde souvent à l'excès. Merci de ne jamais m'avoir reproché de rester plus longtemps avec des amis ou mes parents en France, de toujours avoir compris les opportunités que j'ai pu avoir durant mon doctorat (JUAS, conférences . . .). Je n'oublierais jamais le fait que tu sois venue avec moi ici au Québec, sans te poser de question, le plus beau geste que l'on puisse faire à la personne que l'on aime. Ça nous a permis de devenir addicts aux bagels et au sirop d'érable sous toutes ses formes. . .

À mes parents ! La mif, le sang de la veine comme on dit dans le sud. Merci de m'avoir toujours laissé faire ce que je voulais et de m'accorder votre confiance dans tous mes choix de vie. Être à Pignan c'est toujours un vent d'air frais (Mistral, Tramontane comme on veut), un retour aux sources qui permet de recharger les batteries entre le couscous, les tielles sétoises, la praluline, Pézenas, Saint-Guilhem et j'en passe. Merci infiniment de m'avoir emmené au ping et au basket tant de fois, de m'avoir inculqué cette ouverture d'esprit et des valeurs dont je suis fier d'arborer et pour tous ces voyages ensemble.

À Thibaut, Volpi, Fred et Bi ! Je vous ai tant gonflés avec les stats de Lebron qui soit dit en passant est toujours en 27-7-7 en carrière toujours à 40 ans. Le gars, il va être top 10 de toutes les catégories statistiques ! Merci pour tous ces moments de partage autour d'un match de foot, tennis ou basket, d'une partie de Time Bomb et j'en passe. Thib, je t'ai toujours considéré comme mon grand frère, on a tant partagé de moments précieux depuis 13 ans maintenant. Volpi, travaille ton service monc tu fais trop de doubles fautes ! Fred, en vrai Thierry Henry il a 3 Benzema dans chaque pied ! Bi, t'es rouge monc !

À Robin, Nico, Jerem et Paulo ! Quoi dire les frères. . . Des premières soirées en prépa à ce road trip en camping-car à travers le Québec en passant par les fêtes de Dax ou de Bayonne, j'ai l'impression qu'on a eu 100 vies ensemble. Il nous reste tellement à voir (le mariage de Paul mdr) et de bières à boire !!!

À Axel et Thibault ! Ou devrais-je dire Polochon et Moumoute chef ? Alors quoi dire de ce duo sorti tout droit d'un asile de fous ? Si un jour vous nous avez croisés tous les trois ensemble, vous avez dû avoir les tympans brisés. Entre les discussions philosophiques sur le sens de la vie, celles un peu moins philosophiques dont je tairais le sujet, les parties de gaming sur Rocket League ou BG3 et nos vacances (toujours sous la flotte), on pourrait faire une compile des conneries qu'on a balancées. Cap 33, Miss Gironde, Louloute cœur, la moumoute de Gourdon, le Spar, Datus Locum . . . Bref finalement après 10 ans, une chute de cheveux et 2000 dB par discussion on ne sera toujours pas allé à Hawaï. . .

À Elissa ! Mon exemple depuis que je suis arrivé au Québec et à l'INRS. La directrice et reine de l'INRS. Celle que tout le monde connait et apprécie parce qu'elle est prête à aider tout le monde mais jamais sans sa Pita. J'ai toujours admiré ta capacité à gérer tant de choses à la fois, garder la tête froide quand la pièce autour de toi est feu. Elissa, toi la pression tu la bois même si tu préfères les cocktails avec 2 tonnes de sucre ! Tié la famille ma gâtée !

À Amélie ! Ma cheffe et amie. Tu m'as appris que transmettre est encore plus important qu'apprendre. Se soucier de l'après pour que perdure notre passion est la plus belle des choses. Amélie c'est un train de marchandises lancé à la vitesse d'un TGV, une centrale nucléaire d'idées. Je t'ad-

mire car tu as tant d'initiatives, d'idées innovantes que tu poursuis avec passion. Notre aventure Apprentis-Chercheurs est folle et m'a tant appris sur le fait de transmettre aux générations futures. Merci à toi, le sang !

À ceux qui ne me connaissent pas mais qui font partie de mon quotidien :

- Bastien Fontanieu et Alexandre Martin (Trashtalk Production)

- Marie Beljean, Benoit Maylin, Julien Pichené, Service Volée, Julien Varlet, Rodolphe Cazejuste et Frédéric Verdier (Sans Filet)

- Walid Acherchour, Thomas Bonnavent, Tidiany M'Bo, Elton Mokolo, Romain Beddouk, Najim Medini et Yanis Yemmi (WFC)

- Shai Mamou, Antoine Pimmel et Théophile Haumesser (Reverse Basket Session)

Vous êtes la représentation de ma personnalité et je ne pouvais pas oublier de vous citer car finalement je vous vois et vous entends quasiment tous les jours !

Et puis comment ne pas citer mon idole d'enfance, l'homme aux 14 coupes des mousquetaires et 22 grands chelems : Rafael Nadal ! 24 ans de carrière professionnelle et toujours la même ligne de conduite : le bandeau bien serré et une abnégation sans faille. C'est ce que j'ai le plus admiré chez toi, cette capacité à jouer chaque point comme si ta vie en dépendait mais aussi à t'élever quand il fallait et face au meilleur. J'ai toujours voulu m'inspirer de toi, surtout dans les moments de doute et d'échec. Tu m'as aidé à me dépasser durant mes examens, à me dire que ce n'est pas si grave et que la fois suivante ça ira mieux. Essayer, se remettre en question et repartir de l'avant mais toujours à 200%, c'est ta philosophie et je tente d'en faire la mienne depuis ce mois de mai 2006 durant lequel j'ai commencé à t'admirer.

A special thank you to Canan, my brilliant lab partner and conference partner-in-crime at LPAW 2025. From insightful discussions to our favorite pastime; quietly judging everything and everyone. We made even the most technical sessions unforgettable. Thank you for the laughs and shared sense of mischief throughout this journey.

Il m'est essentiel de remercier chaleureusement mon directeur de thèse, Patrizio Antici. Tu as été un des piliers de mon apprentissage du monde de la recherche, influençant non seulement ma manière de travailler et de voir les choses, mais aussi mon développement personnel en me présentant au monde de la recherche scientifique, tu m'as aidé à devenir une meilleure personne, et pour cela, je vous en suis profondément reconnaissant. Patrizio, depuis cette première rencontre conviviale dans ce bistrot proche de Montparnasse le temps est passé si rapidement ! Ce jour-là, nous avions monté les premières pièces de ce projet long de 3 années et demie. J'ai découvert une personne au courant de tout, omniprésente et engagée à l'échelle internationale. Comment tu fais pour gérer tout ça ? Je ne l'ai jamais compris, mais cela m'a toujours subjugué. Tu m'as également inculqué que tout part d'une bonne manip. Merci, sincèrement.

À mes deux acolytes de Varennes : Sylvain et Simon ! Ce que je sais faire aujourd'hui je l'ai en très grande majorité appris grâce à vous et à votre immense envie de vous améliorer et votre érudition. Sylvain, ta plus grande qualité est celle de ne jamais blâmer qui que ce soit pour ses erreurs, tu cherches avant tout la solution plutôt que la cause de l'erreur. C'est cette vision et ce flegme que j'admire le plus chez toi. Et puis, tu m'as tant appris, aidé, conseillé et suivi sans rien demander en retour dès le début. J'ai énormément appris à tes côtés durant les nombreuses heures qu'on a passées au labo à aligner, calibrer, mesurer, déboguer ... Mais c'était le fun de ouf ! Simon, tu as été la rampe de relancement de mon envie d'entreprendre ce doctorat. Tout a commencé lors de notre première rencontre virtuelle durant laquelle tu étais si enthousiaste de ce que tu avais accompli pendant ta thèse. Cela m'a permis de faire la même chose, d'avoir confiance en ce que je voulais faire et d'en être fier. Ta rigueur est ton plus grand atout, c'est finalement pour cela qu'on est tous les 3 complémentaires. Merci de m'avoir épaulé, d'avoir fait de moi ton padawan.

La route est encore longue mais finalement ce n'est pas la destination qui est importante, ce sont les épreuves qui nous y ont menés avec qui on les a surmontées !

À Joël le djeun's ! Joël c'est un cap que dis-je, c'est une péninsule ! Si vous lui demandez si "C'est une bonne situation, ça, chercheur ?", il vous répondra qu'il n'y a pas de bonne ou mauvaise situation mais que s'il devait résumer sa vie aujourd'hui ce sont surtout les rencontres. Tu as été d'une aide précieuse dans ma thèse, "et ça, ça fait plaisir !". Grâce à toi, j'ai pu avancer "tranquille, peinard, à la fraîche", sans trop de sueur (enfin, presque). Quand j'avais des problèmes, "tu es venu, tu as vu, et tu as vaincu", avec une efficacité redoutable. Honnêtement, "ce que tu fais, c'est bien, continue", parce que sans toi, j'aurais fini "à genoux dans le sable en train de pleurer". Alors, un grand merci, "avec les mains, avec le cœur, avec le foie", et surtout avec "un peu de potion magique" pour fêter ça !

I would also like to express my heartfelt gratitude to LaserNetUS for the distinct honor of being named one of the first two LaserNetUS Student Ambassadors. This recognition has provided a unique opportunity to engage meaningfully with the high-intensity laser community, contribute to collaborative research initiatives, and advance my work in laser-driven ion acceleration. To me, LaserNetUS is a unique network of high-power laser facilities across North America, funded by the U.S. Department of Energy, with the mission of promoting broad access to cutting-edge laser infrastructure and fostering collaborative research in high-intensity laser science. Through workshops, facility visits, and discussions with leading scientists across institutions, I was able to broaden my scientific perspective while advancing my research in laser-driven ion acceleration. The recognition and responsibilities associated with this ambassadorship have been both an honor and a catalyst for my growth as a researcher. I am deeply grateful to the LaserNetUS consortium for their support, vision, and commitment to training the next generation of scientists in this rapidly evolving field.

RÉSUMÉ

Les accélérateurs d'ions générés par interaction laser-matière représentent une voie prometteuse pour de nombreuses applications en science des matériaux, en médecine, en radiographie ultrarapide et en physique nucléaire. En effet, les lasers femtosecondes de puissance ultra-élevée, focalisés à des intensités extrêmes, permettent de produire des faisceaux d'ions compacts atteignant plusieurs mégaélectronvolts. Parmi les mécanismes d'accélération disponibles, l'accélération TNSA (« Target Normal Sheath Acceleration ») demeure aujourd'hui la plus robuste et la plus utilisée. Toutefois, la stabilité, la qualité et l'efficacité énergétique des faisceaux générés restent des enjeux majeurs. En particulier, les performances résultent d'un équilibre subtil entre des mécanismes de couplage laser–plasma hautement non linéaires, sensibles à la géométrie du faisceau, aux propriétés de la cible et à la structure spatio-temporelle de l'impulsion laser. Les progrès expérimentaux sont souvent ralentis par les faibles cadences de tir, le bruit de mesure et le coût élevé des campagnes expérimentales nécessitant l'exploration de larges espaces de paramètres. Bien que les simulations numériques puissent offrir un éclairage précieux, les modèles réalistes demeurent coûteux en temps de calcul et ne remplacent pas une optimisation expérimentale systématique.

Parallèlement, l'essor de l'intelligence artificielle et des stratégies de contrôle autonome ouvre de nouvelles perspectives pour les accélérateurs laser–plasma. En particulier, l'optimisation bayésienne constitue un outil efficace pour explorer rapidement des espaces de paramètres complexes tout en réduisant le nombre d'expériences nécessaires. Ces approches permettent d'accélérer l'optimisation, d'améliorer la reproductibilité et de réduire la dépendance aux ajustements manuels. Cette thèse présente la mise en place d'une plateforme d'accélération d'ions stable et partiellement automatisée hébergée par l'« *Advanced Laser Light Source* (*ALLS*)» de l'Institut national de la recherche scientifique (INRS). Ce travail combine développement expérimental, modélisation numérique et optimisation guidée par apprentissage automatique afin d'améliorer les performances et la stabilité des faisceaux d'ions. Après avoir renforcé l'alignement, la préparation des cibles, l'intégration des diagnostics et la répétabilité du système, nous avons conçu une méthode d'optimisation multi-étapes fondée sur des modèles de substitution et l'optimisation bayésienne, adaptée aux contraintes des lasers à haute puissance. Cette approche a permis d'optimiser efficacement plusieurs paramètres expérimentaux malgré le bruit et le faible taux de répétition, conduisant à des améliorations significatives de la stabilité et des performances des faisceaux d'ions. Ces travaux contribuent à établir l'accélération laser-plasma assistée par apprentissage automatique comme une voie prometteuse pour les futures sources d'ions compactes et leurs applications scientifiques, industrielles et sociétales.

Mots-clés : Interaction laser-plasma; Accélérateurs de protons par laser; Target Normal Sheath Acceleration; Automatisation; Apprentissage automatique; Optimisation bayésienne; Modelage du front d'onde laser; Fluctuations tir à tir; Stabilité des faisceaux d'ions; Contrôle adaptatif des faisceaux d'ions.

ABSTRACT

Laser-driven ion accelerators are promising candidates for applications in materials science, medical therapy, ultrafast probing, and nuclear physics. Indeed, modern high-intensity laser systems, operating in the multi-hundred-terawatt range and focused to extreme intensities, can generate multi-MeV proton beams from compact, high-brightness sources. Among the available acceleration schemes, Target Normal Sheath Acceleration (TNSA) remains the most robust and widely used, but improving the stability, beam quality, and energy throughput of these sources remains a central challenge. In particular, their performance results from a complex interplay of nonlinear laser–plasma processes that are highly sensitive to laser conditions, target properties, and spatio-temporal laser beam structure. Furthermore, progress is often limited by low repetition rates, substantial experimental noise, and the cost of large multi-parameter scans. Although numerical simulations can provide valuable insight, high-fidelity models remain computationally expensive and cannot replace systematic experimental optimization.

In parallel, the rapid development of data-driven methods and autonomous control strategies has opened a new path for laser-plasma accelerators. Machine learning, and in particular Bayesian optimization, offers a powerful route to efficiently explore high-dimensional parameter spaces while requiring only a limited number of experimental evaluations. Its successful deployment in other accelerator platforms demonstrates its potential to accelerate optimization, reduce human intervention, and improve overall machine reliability.

This thesis presents the development of a stable and partially automated laser-driven ion acceleration platform at the *Advanced Laser Light Source* (*ALLS*) facility of the *Institut national de la recherche scientifique* (*INRS*). The work combines experimental innovation, numerical modeling, and machine-learning-assisted control to enhance both performance and stability. After improving target alignment, targetry, diagnostic integration, and operational repeatability, we developed a multi-step surrogate-assisted Bayesian optimization framework tailored to the constraints of high-power laser systems. This approach enabled efficient optimization of key experimental parameters despite low repetition rates and significant shot-to-shot fluctuations. The resulting improvements in beam stability and proton performance demonstrate that data-driven strategies can effectively guide experiments, paving the way toward adaptive, autonomous laser-driven ion accelerators. This work contributes to establishing machine-learning-assisted laser-plasma acceleration as a promising route for future compact ion sources and their scientific, industrial, and societal applications.

Keywords : Laser-plasma interaction; Laser-driven proton accelerators; Target Normal Sheath Acceleration; Automation; Machine Learning; Bayesian optimization; Laser wavefront shaping; Shot-to-shot fluctuations; Ion beam stability; Adaptive ion beam control.

SOMMAIRE RÉCAPITULATIF : OPTIMISATION PAR APPRENTISSAGE AUTOMATIQUE DES LIGNES DE FAISCEAUX D'IONS ACCÉLÉRÉS PAR INTERACTION LASER-MATIÈRE

Contexte général et hypothèse directrice

Cette thèse s'inscrit dans le développement d'accélérateurs d'ions générés par interaction laser-matière dans le régime *Target Normal Sheath Acceleration* (TNSA) (Wilks et al., 2001) et, plus généralement, des accélérateurs laser-plasma. Si des protons de plusieurs MeV peuvent être accélérés de manière routinière via TNSA (Zeil et al., 2010; Fourmaux et al., 2013; Yao et al., 2024b), l'adoption de ces sources pour des usages scientifiques, technologiques ou sociétaux (Bolton et al., 2018) exige des améliorations au niveau de leur stabilité, répétabilité et efficacité. L'hypothèse centrale de cette thèse est qu'une optimisation assistée par apprentissage automatique (*Machine Learning*, ML) — en particulier l'optimisation bayésienne (OB) (Jones et al., 1998; Brochu et al., 2010; Shahriari et al., 2015; Frazier, 2018) — peut améliorer de façon significative la performance et la robustesse des faisceaux d'ions accélérés, y compris dans des conditions expérimentales à faible cadence et marquées par un bruit intrinsèque élevé, caractéristiques du domaine de l'accélération laser-plasma (Nishiuchi et al., 2009; Treffert et al., 2022a; Giuffrida et al., 2023; Nedbailo et al., 2023). Cette hypothèse est testée sur la ligne d'ions de l'*Advanced Laser Light Source* (*ALLS*) (Vallières et al., 2020), installée au Centre Énergie Matériaux Télécommunications de l'Institut national de la recherche scientifique (INRS) dans la province du Québec au Canada, en exploitant un contrôle spatio-temporel du faisceau laser, utilisé pour la génération et l'accélération des ions, à l'aide d'un miroir déformable (MD) et d'un filtre dispersif programmable (FDP).

Objectifs

Les travaux présentés dans cette thèse poursuivent trois objectifs principaux.

Tout d'abord, Il s'agit d'abord de démontrer qu'une source d'ions accélérés par le mécanisme TNSA peut fonctionner de manière stable sur un grand nombre de tirs consécutifs, tout en caractérisant avec précision ses grandeurs caractéristiques. Celles-ci incluent notamment l'énergie maximale des protons ($\mathcal{E}_{p,max}$), les flux d'ions différentiels ($\frac{d^2N}{d\mathcal{E}d\Omega}$) et totaux ($N_{tot}$), ainsi que l'efficacité de conversion de l'énergie laser en énergie des ions accélérés (α).

Le second objectif est d'évaluer et de comparer l'influence relative de deux leviers de contrôle indépendants sur les grandeurs caractéristiques des faisceaux de protons : d'une part, l'ajustement du profil temporel de l'impulsion laser via sa phase spectrale, en modulant la dispersion de groupe (GDD) et la dispersion d'ordre trois (TOD) ; d'autre part, le profilage spatial du front d'onde du faisceau au moyen d'un MD. Cette analyse vise à déterminer lequel de ces paramètres, ou quelle combinaison, contribue le plus efficacement à l'amélioration des performances du mécanisme de TNSA.

Enfin, le troisième objectif est de concevoir et valider une stratégie d'optimisation multi-étapes (*Multi-Step Bayesian Optimization*, MSBO) intégrant des modèles de ML de type forêts aléatoires

et une OB. L'enjeu est de démontrer expérimentalement que cette approche permet d'identifier rapidement et efficacement une configuration optimale, puis d'en quantifier le gain en performance par rapport aux méthodes d'optimisation conventionnelles.

Approche expérimentale et dispositifs

Les expériences ont été menées sur la ligne d'ions de l'installation *ALLS*, sous différentes configurations expérimentales (Figures 3.1.A, 4.6 et 5.1), en utilisant des impulsions laser polarisées p, incidentes à 20° sur la cible et tirant à une cadence de 0.625 Hz. Les cibles consistent en des feuilles métalliques d'épaisseur micrométrique, adaptées à l'accélération d'ions en régime TNSA.
La caractérisation des faisceaux est assurée à l'aide d'un spectromètre à parabole de Thomson (TPS) équipé d'un détecteur à microcanaux (MCP), permettant de mesurer avec précision les spectres d'énergie et les grandeurs caractéristiques des ions accélérées. Le contrôle spatio-temporel du faisceau laser est assuré par l'utilisation combinée d'un MD couplé à un analyseur de front d'onde et d'un FDP.
Afin de permettre l'acquisition de données sur de longues séries expérimentales, un porte-cibles multi-échantillons, d'une capacité de 400 cibles, a été conçu et intégré au système expérimental (voir Figure 3.1.C).

Développement

Chapitre 1 - Fondements physiques et technologiques de l'accélération laser-plasma

Ce chapitre présente les fondements technologiques et physiques de l'accélération laser–plasma, en s'articulant autour de : (i) l'essor des lasers ultra-intenses permis par la *Chirped Pulse Amplification* (CPA) (Strickland et al., 1985), (ii) le cadre électromagnétique, fondé sur les équations de Maxwell, régissant l'interaction laser–plasma, et (iii) les principaux mécanismes d'accélération d'électrons et d'ions, avec un accent particulier sur le régime TNSA. L'introduction rappelle le principe en trois étapes de la CPA — étirement, amplification et recompression — (Figure 1.1), puis replace ces avancées dans le contexte international des grandes installations (Figure 1.2) et de l'évolution historique des puissances crêtes et intensités atteintes (Figure 1.3). La montée en cadence et en puissance moyenne, essentielle pour les plateformes expérimentales modernes, est illustrée par la cartographie énergie–cadence de la Figure 1.4.

Le cadre électromagnétique est rappelé sous forme différentielle pour traiter les phénomènes ultrarapides et localisés caractéristiques des interactions laser–plasma (Maxwell–Lorentz). Dans les plasmas sous-denses, l'accélération d'électrons par sillage (*Laser Wakefield Acceleration*, LWFA) met en jeu la force pondéromotrice (Esarey et al., 2009) et la formation d'une cavité ionique («bubble», Figure 1.5), permettant d'atteindre des gradients de l'ordre de 100 GV/m grâce à des cibles gazeuses de densité électronique $n_e \sim 10^{18}\,\mathrm{cm}^{-3}$ et la production concomitante de rayonnement X betatron (Esarey et al., 2002; Kostyukov et al., 2003) (Figure 1.6). Ces résultats motivent, côté ions, l'étude des régimes à cible solide où dominent les champs quasi-électrostatiques de gaine.

Le mécanisme TNSA, mis en évidence expérimentalement dès les années 2000 (Clark et al., 2000; Maksimchuk et al., 2000; Snavely et al., 2000; Wilks et al., 2001) et décrit analytiquement via des modèles d'expansion plasma dans le vide (Mora, 2003), repose sur la génération d'élec-

trons chauds à la surface frontale de la cible, leur transport balistique à travers celle-ci, puis la formation d'une gaine électrostatique à l'arrière, responsable de l'accélération des protons (Figure 1.7) et d'autres espèces ioniques. Le déroulement temporel de ces étapes clefs est synthétisé dans la Figure 1.8. Les mécanismes d'absorption conduisant à la production d'électrons chauds dépendent fortement de l'intensité laser, de la polarisation et de l'angle d'incidence. On distingue principalement : l'absorption résonante (Wilks et al., 1992) (Figure 1.9), le chauffage de Brunel (Brunel, 1987) en présence de gradients abrupts de densité, et, au-delà du seuil relativiste, le chauffage $\vec{J} \times \vec{B}$ (Kruer et al., 1985), qui domine alors la dynamique et fixe l'ordre de grandeur de la température électronique via l'énergie pondéromotrice. Le phénomène de saturation de l'absorption est discuté en lien avec la Figure 1.10. Dans un second temps, le modèle isotherme à durée d'accélération finie relie la température des électrons chauds T_{hot} (Mora, 2003; Fuchs et al., 2006), la géométrie de la gaine et l'échelle temporelle d'injection à l'énergie maximale des ions. Enfin, les principales lois d'échelle établies à partir d'études paramétriques et d'ensembles expérimentaux (Fuchs et al., 2006; Robson et al., 2007; Flippo et al., 2008; Zeil et al., 2010) sont synthétisées et mises en perspective avec la Figure 1.11.

Au-delà du régime TNSA, d'autres mécanismes prennent le relais lorsque l'intensité atteint des valeurs extrêmes ($\geq 10^{20}$ W.cm^{-2}) et/ou lorsque la cible devient ultrafine (épaisseurs sub-micrométriques). Parmi eux, la *Radiation Pressure Acceleration* (RPA) (Esirkepov et al., 2004) constitue le scénario radiatif de référence : la pression de radiation transmet directement sa quantité de mouvement au plasma, impulsionnant collectivement les ions. Ce régime se décline en deux configurations — *Hole Boring* (HB) pour des cibles « épaisses » et *Light Sail* (LS) (Qiao et al., 2009; Macchi et al., 2013) pour des cibles nanométriques — comme illustré en Figure 1.12. La polarisation circulaire, en réprimant le chauffage collisionless de type $\vec{J} \times \vec{B}$, renforce l'efficacité de cette poussée collective. À encore plus haute intensité, l'augmentation de la température électronique peut conduire à une transition vers la transparence relativiste (Vshivkov et al., 1998), ouvrant la voie à des scénarios d'accélération volumique tels que le *Break-Out Afterburner* (BOA) (Yin et al., 2006, 2007). Dans ces conditions, l'onde laser pénètre le plasma devenu transparent au sein relativiste du terme, ce qui favorise un chauffage volumique efficace et l'établissement de champs longitudinaux capables de générer des spectres ioniques plus énergétiques. D'autres mécanismes, tels que la *Collisionless Shock Acceleration* (CSA) (Silva et al., 2004; Haberberger et al., 2012) et le *Magnetic Vortex Acceleration* (MVA) (Nakamura et al., 2010; Park et al., 2019), sont également rapportés dans la littérature pour certaines configurations de densité et de durée d'impulsion. Bien qu'ils ne soient pas étudiés en détail dans ce manuscrit, ils sont mentionnés pour compléter le panorama des moteurs d'accélération laser-plasma accessibles à haute intensité.

Les performances et la stabilité des faisceaux accélérés dépendent à la fois des caractéristiques du laser (énergie, durée d'impulsion, contraste temporel, polarisation, géométrie focale) et des propriétés des cibles (matériau, épaisseur, densité, structuration). Pour répondre aux contraintes imposées par les cadences élevées, plusieurs solutions de rafraîchissement ou de renouvellement des cibles ont été développées, telles que les rubans (Noaman ul Haq et al., 2017; Noaman ul Haq et al., 2018; Dover et al., 2020b), les jets liquides (Gauthier et al., 2016; Obst-Huebl et al., 2017; Grieser et al., 2019; Polz et al., 2019), les systèmes multi-échantillons (Catrix et al., 2023), ou encore les jets de gaz (Jinno et al., 2013; Chen et al., 2017; Puyuelo-Valdes et al., 2019b), illustrés en Figure 1.13. Du côté de la métrologie, la caractérisation spectrale et multi-espèces repose principalement sur les détecteurs à parabole de Thomson (Sadowski et al., 2009; Gwynne et al., 2014; Tata et al., 2017), dont les schémas et l'implantation expérimentale sont présentés en Figure 1.14. Ce diagnostic est couramment complété par des mesures de charge à l'aide d'*Integrating*

Current Transformers (ICT) (Geulig et al., 2022), par des réseaux de fibres scintillantes (Patel et al., 2024), ainsi que par des dispositifs temps-de-vol (Scuderi et al., 2020; Vallières et al., 2020).

Enfin, le chapitre replace ces sources dans leurs domaines applicatifs : analyses élémentaires et science des matériaux (Barberio et al., 2017b) (*Particle Induced X-ray Emission* et imagerie, Figures 1.15.A–C), chauffage isochore (Patel et al., 2003; Roth et al., 2005) (Figure 1.15.D), radiobiologie/hadronthérapie et scénarios FLASH (Vozenin et al., 2022; Flacco et al., 2024), radiographie protonique de champs transitoires (Mackinnon et al., 2004), et production radio-isotopique médicale (Ledingham et al., 2004) (Figure 1.15.E). L'ensemble établit le lien entre technologies CPA, cadres électromagnétiques et stratégies expérimentales, préparant les méthodes d'implantation, de contrôle et d'optimisation développées dans les chapitres suivants.

Chapitre 2 — *Machine Learning* dans la physique des plasmas

Ce chapitre introduit les concepts fondamentaux de ML et examine son rôle émergent dans la physique des plasmas, en particulier pour l'optimisation de l'accélération d'ions par laser (Hatfield et al., 2021; Trieschmann et al., 2023; Döpp et al., 2023; Roussel et al., 2024). Après avoir défini le ML et ses principales sous catégories — apprentissage supervisé, non supervisé et par renforcement — (Figures 2.1 et 2.2), il présente son fonctionnement typique (Figure 2.3) et ses limites actuelles.

L'intégration du ML en physique des plasmas répond à la complexité non linéaire et multidimensionelle des interactions laser–plasma, aussi bien en régimes de plasmas froids que chauds, comme le montre la comparaison synthétique du Tableau 2.1. Les applications incluent la diminution du coût computationnel de simulations, la modélisation de plasmas complexes et l'optimisation de procédés expérimentaux.

La section suivante analyse la complexité des interactions laser–plasma : chauffage électronique, auto-focalisation relativiste (Ritchie, 1994; Sprangle et al., 2007) ou pondéromotrice (Osman et al., 1999) (Figure 2.4), instabilités paramétriques (Giulietti et al., 1999; Salcedo et al., 2003; Yao et al., 2024a; Ruskov et al., 2024; Lian et al., 2025), et importance des diagnostics expérimentaux comme la diffusion Thomson (Kaloyan et al., 2021, 2022) ou la radiographie protonique (Mackinnon et al., 2004; Romagnani et al., 2005).

Le chapitre détaille ensuite la nature multi-paramétrique de l'accélération d'ions par laser, influencée par :

- les paramètres laser (énergie, durée, contraste temporel, qualité du front d'onde),
- les propriétés de la cible et les conditions de pré-plasma (Hadjisolomou et al., 2020; Gizzi et al., 2021; Kim et al., 2022; Strehlow et al., 2022; Keppler et al., 2022),
- l'environnement expérimental (pression résiduelle (Snyder et al., 2020), champs magnétiques externes (Sarri et al., 2012; Ferri et al., 2019a; Weichman et al., 2020; Khan et al., 2024)).

Ces paramètres doivent être optimisés conjointement pour maximiser l'efficacité et la stabilité du faisceau d'ions.

Différentes techniques ML appliquées à l'accélération laser-plasma sont ensuite évoquées :

— **Réseaux de neurones** : puissants pour la modélisation non linéaire et la prédiction de spectres de protons (Figure 2.5);

— **Support Vector Regression** (SVR) : adaptée aux jeux de données synthétiques et non bruités (Figure 2.6) ;

— **Forêts aléatoires** : modélisent efficacement des relations non linéaires complexes tout en restant robustes au bruit expérimental et interprétables malgré des jeux de données limités;

— **Régression par processus gaussiens** (GPR) : fournit à la fois prédictions et incertitudes (Figure 2.8) ;

— **Optimisation bayésienne** (OB) : exploite généralement la régression par processus gaussiens pour optimiser des fonctions coûteuses à évaluer, en équilibrant exploration et exploitation (Figure 2.9);

— **Apprentissage par renforcement** (RL) : contrôle adaptatif en temps réel d'installations complexes (Figure 2.10).

Chaque méthode est comparée en termes de domaines d'efficacité, points forts et limites (Tableau 2.6).

En conclusion, le chapitre souligne l'impact du ML sur l'efficacité expérimentale (réduction du temps d'optimisation, diagnostic synthétique en temps réel, contrôle adaptatif) et sur la compréhension des interactions laser–plasma (découverte de corrélations cachées, lois d'échelle révisées, optimisation multi-objectifs). Les perspectives incluent : approches hybrides (GPR + RL), réseaux de neurones informés par la physique sous-jacente des interactions laser-plasma, intégration de diagnostics synthétiques en boucle fermée et utilisation des hautes performances de calcul des machines futures (ordinateurs quantiques, supercalculateurs...) pour accélérer les simulations et l'optimisation expérimentale.

Chapitre 3 — Fonctionnement répété et stabilité

Ce chapitre présente la mise en œuvre et la caractérisation de la ligne d'accélération d'ions d'*ALLS*, optimisée pour le régime TNSA (Catrix et al., 2023). Le système repose sur un porte-cibles motorisé multi-cibles (jusqu'à 400 cibles, Figure 3.1.C), permettant l'acquisition de longues séries de tirs consécutifs avec un maintien de l'alignement par un système de positionnement à cinq axes (Figure 3.4). L'importance de ce type de montage réside dans la capacité à enchaîner des centaines de tirs sans réalignement manuel, ce qui était auparavant l'un des principaux obstacles à la reproductibilité des expériences en accélération laser-plasma. La motorisation à plusieurs degrés de liberté permet de compenser les dérives mécaniques et d'alignement optique, assurant une meilleure reproductibilité dans le temps. Les faisceaux d'ions produits par TNSA sont caractérisés à l'aide d'un TPS couplé à un MCP calibré (Figure 3.5, Tableau 3.1), et les spectres sont analysés par un traitement automatisé via *Matlab* pour extraire les grandeurs caractéristiques des faisceaux d'ions. L'automatisation de la chaîne de traitement, depuis l'acquisition des images jusqu'au calcul des spectres et des grandeurs caractéristiques des faisceaux d'ions, constitue un gain en fiabilité et en rapidité, et permet également de manipuler un volume de données beaucoup plus important qu'avec des analyses manuelles, ouvrant la voie à des études statistiques robustes sur plusieurs centaines de tirs consécutifs.

Les mesures effectuées sur une série de 296 tirs consécutifs sur une cible d'aluminium de 4.5 µm d'épaisseur (Figures 3.6.A–B) montrent une excellente stabilité au niveau de l'énergie maximale

des protons : $\mathcal{E}_{p,max} = (5.6 \pm 0.4)$ MeV, soit $\sim$7 % de fluctuation, et des variations de flux intégré (N_{tot}) de 14.9 %. Cette longue série démontre que l'on peut maintenir un fonctionnement stable sur plusieurs minutes ($\sim$15 minutes) sans réalignement, alors même que chaque tir détruit localement la cible. Cette performance est directement liée à la rigidité du porte-cibles et à la stabilité du laser, qui ne présente que 2.5 % RMS de variation d'énergie impulsionnelle. Ce niveau de stabilité se démarque par rapport aux autres infrastructures internationales, où les fluctuations rapportées dépassent souvent 20–30 % pour des systèmes similaires. L'efficacité de conversion entre l'énergie du laser et l'énergie des protons reste typique de ce type d'infrastructure $\sim$0.005 avec une fluctuation tir à tir de 12 %. Ces valeurs de fluctuations tir à tir, bien que modestes en comparaison des grands accélérateurs conventionnels, sont cohérentes avec les mécanismes TNSA où seule une fraction des électrons chauds contribue effectivement à la formation du champ de gaine. Elles témoignent néanmoins d'une remarquable constance sur la durée de l'expérience. Les comparaisons multi-matériaux et multi-épaisseurs (Figures 3.6.C et 3.7, Tableau 3.3) confirment que :

— les cibles métalliques fines (2 µm Ta, 1 µm Cu) produisent les énergies maximales les plus élevées (~ 6.2 MeV) avec une bonne stabilité (5-7 %),

— l'aluminium 4.5 µm offre la meilleure robustesse mécanique pour les longues séries, malgré un rendement moindre,

— les cibles en Polytéréphtalate d'éthylène (PET) présentent des fluctuations importantes (~ 45 %), dues à leur souplesse mécanique et à la difficulté d'alignement optique. Leur transparence aux diagnostics optiques rend de surcroît le centrage plus complexe, expliquant leur instabilité par rapport aux métaux opaques et rigides.

Ces résultats soulignent que la reproductibilité des faisceaux ne dépend pas uniquement du laser mais de la combinaison fine entre matériau, épaisseur et contraintes mécaniques, ce qui impose une réflexion sur le design global du système.

À intensité réduite ($I_L \sim 3.3 \times 10^{19}$ W.cm^{-2}), les énergies maximales chutent à ~ 1.8 MeV (Figure 3.9.A), mais la stabilité relative d'$\mathcal{E}_{p,\max}$ reste comparable (5–9 %). Cela prouve que, même dans un régime où l'énergie déposée est plus faible et où les électrons produits ont une température moindre, la reproductibilité au niveau de l'énergie maximale du système n'est pas dégradée. En revanche, le flux à 1 MeV de protons devient plus fluctuant (Figure 3.9.B), en raison d'une plus grande sensibilité aux irrégularités de surface et aux variations d'absorption à plus faible intensité laser. Ces observations confirment l'importance d'un excellent contraste temporel : en l'absence de préplasma stabilisant, la rugosité locale ou les défauts de surface dominent la physique de l'absorption, rendant les fluctuations plus marquées.

Enfin, l'application à la spectroscopie PIXE/XRF (*Particle Induced X-ray Emission/X-ray Fluorescence*), illustre concrètement l'intérêt d'une source TNSA versatile et stable à haut taux de répétition. Dans cette configuration, un faisceau d'ions et de rayons X généré par une cible de Ta d'épaisseur 2 µm est utilisé pour exciter un échantillon organique contenant des éléments traces. L'accumulation de 50 tirs consécutifs, rendue possible grâce à la stabilité mécanique et énergétique du système, permet de faire émerger dans le spectre des raies d'émission caractéristiques d'éléments de faible concentration, en particulier l'arsenic (As) avec une raie K_α détectée à 10.54 keV (Figure 3.10). En comparaison, un tir unique ne permet pas de distinguer cette raie du bruit de fond, démontrant que la répétabilité et la constance du faisceau sont essentielles pour atteindre des seuils de détection plus bas. Cette capacité à accumuler des tirs reproductibles ouvre la voie à des applications analytiques avancées dans le domaine du patrimoine culturel, de la science des matériaux ou de la médecine, où la détection de traces élémentaires peut avoir

une valeur critique. En particulier, elle montre que l'accélération laser-plasma n'est pas seulement une prouesse physique, mais qu'elle peut devenir un véritable outil d'analyse multidisciplinaire.

Ce résultat met en évidence que l'optimisation conjointe du matériau et de l'épaisseur des cibles ne répond pas uniquement à des critères de rendement énergétique ou de stabilité, mais constitue également un levier pour adapter la source de protons aux besoins spécifiques de différentes applications avancées. Parmi celles-ci figurent la caractérisation d'éléments traces, l'imagerie résolue en énergie et les études de composition élémentaire de matériaux sensibles, où la possibilité de cumuler des acquisitions sans dégradation significative des performances est un atout déterminant. En outre, ce travail démontre l'importance des approches comparatives et statistiques, car ce n'est que sur de longues séries de tirs que la notion de stabilité prend tout son sens. Les résultats obtenus à l'INRS placent ainsi cette ligne d'accélération parmi les systèmes les plus stables actuellement rapportés dans la littérature. En somme, ce chapitre démontre que la mise en place d'une infrastructure robuste et reproductible n'est pas seulement une prouesse technique, mais une condition essentielle pour transformer l'accélération laser-plasma en un outil fiable au service d'applications scientifiques et technologiques concrètes, allant de la physique fondamentale jusqu'aux analyses appliquées.

Chapitre 4 — Contrôle paramétrique des faisceaux de protons par modelage spatio-temporelle du laser

Ce chapitre explore de manière systématique l'influence du modelage spatio-temporel du faisceau laser sur les performances de l'accélération de protons dans le régime TNSA. L'approche expérimentale repose sur le couplage d'un MD bimorphe à 48 actionneurs (Figure 4.7) avec un analyseur de front d'onde et d'un FDP, permettant un contrôle indépendant du front d'onde spatial et de la phase spectrale. Concrètement, l'impulsion Ti :Saphir (3.2 J sur cible, 22 fs minimum, 800 nm) est comprimée par une chaîne CPA double, avec une gestion de phase spectrale par un FDP de type *Dazzler*; le front d'onde est mesuré par un analyseur de front d'onde de type *SID4* et corrigé par un MD bimorphe en boucle fermée. Les mesures sont réalisées sur des cibles d'aluminium de 4.5 µm, irradiées à une intensité de $\sim$1.3$\times$10^{20} W.cm^{-2} ou inférieure, avec un diagnostic de type parabole de Thomson couplé à un MCP (Figure 4.6). Le faisceau est focalisé par une $f/3$ OAP, en polarisation p à 20° de l'axe normal de la cible; la cadence expérimentale est de 0.625 Hz avec un porte-cibles multi-cibles ($\leq$400 cibles) assurant la répétabilité. Un étage XPW en amont améliore le contraste (ASE $\leq$ 10^{-10} à -100 ps). L'extraction des observables ($\mathcal{E}_{p,max}$, N_{tot}, $\mathcal{E}_{tot}$, $\bar{\mathcal{E}}$, $\bar{\mathcal{E}^2}$, T_{hot}) est réalisée tir à tir par un code *Matlab*.

L'étude du modelage temporel, par variation systématique de la dispersion d'ordre 2 (β_2) et de la dispersion d'ordre 3 (β_3) dans la plage $\beta_2 \in [-1250, 1250]$ fs^2 et $\beta_3 \in [-7000, 7000]$ fs^3, montre que les fluctuations des métriques de faisceau de protons restent dominées par le bruit intrinsèque tir-à-tir (Figure 4.8, Tableau 4.1). Pour chaque couple (β_2, β_3), trois tirs consécutifs sont acquis (N=114 configurations au total), ce qui permet d'estimer la variabilité tir-à-tir σ_{sts} et la dispersion inter-configuration σ_{config}; on observe systématiquement $\sigma_{\mathrm{config}} \lesssim \sigma_{\mathrm{sts}}$, indiquant l'absence d'effet systématique fort du façonnage temporel dans la plage explorée. L'élargissement de la durée d'impulsion de 22 fs jusqu'à $\sim$160 fs (Figure 4.8) induit, à énergie laser constante, une baisse d'intensité crête ($I \propto 1/\tau_L$) sans bénéfice corrélatif sur les performances du faiceau de protons, ce qui demeure cohérent avec les cartes iso-intensité de (Fuchs et al., 2006). Par ailleurs, comme le temps de double-transit électronique ($\sim$32 fs pour une cible de 4.5 µm) n'est pas dépassé par l'impulsion courte (22 fs), le mécanisme de renforcement du champ de gaine par recirculation

(Mackinnon et al., 2002) ne peut s'installer; à l'inverse, l'allongement de la durée d'impulsion en variant β_2 et β_3 dégrade potentiellement le contraste et favorise une expansion prématurée, neutralisant tout gain lié au temps d'interaction additionnel avec une impulsion laser plus longue dans le temps. En résumé, dans notre fenêtre paramétrique, une compression maximale de l'impulsion laser proche de la limite de diffraction suffit et un ajustement plus fin de la phase spectrale n'apporte qu'un bénéfice marginal.

À l'inverse, le modelage spatial via le MD génère des variations systématiques et reproductibles sur toutes les métriques de faisceau (Tableau 4.2). Une analyse corrélative (Figure 4.10) révèle qu'un petit sous-ensemble d'actionneurs, en particulier l'Actionneur 0 (défocus global), joue un rôle prépondérant avec des corrélations > 0.8 sur $\mathcal{E}_{p,max}$ et $\mathcal{E}_{tot}$. La campagne d'exploration repose sur 360 configurations issues d'un *Latin Hypercube Sampling* dans l'espace $[-80, 80]$ V pour les 20 actionneurs centraux (0–19), superposées au point de fonctionnement *SID4* (front d'onde quasi-plan). Les corrélations inter-actionneurs faibles ($< |0.17|$) suggèrent une quasi-indépendance des degrés de liberté, ce qui renforce l'interprétabilité des effets de chaque actionneur. La dépendance non linéaire $\mathcal{E}_{p,max}(V_0)$ (Figure 4.13, ajustement quadratique avec $a = -9.96 \times 10^{-4}$ MeV.V^{-2}, $b = 1.28 \times 10^{-2}$ MeV.V^{-1}, $c = 5.34$ MeV) traduit le rôle de la courbure imposée par l'Actionneur 0 sur la taille de la tache focale et la distribution d'intensité, modulant directement T_{hot} et la force du champ de gaine accélérateur.

Des cartes bidimensionnelles issues de modèles de forêts aléatoires (Speiser et al., 2019) (Figure 4.14) montrent que les meilleures performances sont obtenues pour des impulsions temporellement comprimées combinées à une légère tension positive sur l'Actionneur 0, minimisant le défocus. Les modèles RF (composés de 300 arbres décisionnels chacun) sont entraînés sur (τ_L, V_0) et prédisent chacune des métriques des faisceaux de protons; ils atteignent des R^2 de 0.91 à 0.94 selon la métrique, pour des RMSE proches du bruit expérimental, ce qui valide leur usage pour prédire les différentes métriques des faisceaux de protons. La corrélation forte entre $\mathcal{E}_{p,max}$ et l'ensemble des autres métriques (Figure 4.12) confirme sa pertinence comme paramètre unique de suivi et d'optimisation du mécanisme de TNSA, grâce également à sa faible sensibilité au bruit ($\sim$5 % de fluctuation). Ainsi, piloter l'expérience sur $\mathcal{E}_{p,max}$ revient, dans notre cas, à co-optimiser simultanément l'ensemble des métriques des faisceaux de protons, tout en limitant l'influence des incertitudes d'étalonnage et des fluctuations tir à tir.

Dans notre étude, nous mettons en évidence que l'absence de gain lors de l'allongement d'impulsion est liée à la baisse d'intensité pic et aux dégradations possibles du contraste, empêchant l'effet bénéfique attendu de la recirculation électronique. Cette lecture est cohérente avec la littérature : l'allongement n'est avantageux que si la durée dépasse le double-transit (densification de gaine), et si le contraste reste suffisant pour éviter la pré-expansion de la face arrière (Mackinnon et al., 2002; Fuchs et al., 2006). À l'inverse, le modelage spatial permet un contrôle déterministe du couplage, indépendamment de l'intensité maximale, et constitue un levier puissant pour améliorer la performance du mécanisme de TNSA. En pratique, cela se traduit par des rapports $\sigma_{\mathrm{config}}/\sigma_{\mathrm{sts}}$ compris entre 2.2 et 3.8 (Tableau 4.2) pour toutes les observables, avec un maximum sur $\mathcal{E}_{p,max}$.

Sur le plan méthodologique, on notera que les contrastes apparents entre Figures 4.10 et 4.14 résultent des conditions d'échantillonnage : la corrélation de Pearson sous-estime l'influence d'un actionneur lorsque la réponse en énergie est non monotone (symétrie autour d'un optimum de V_0), car des contributions positives et négatives au numérateur se compensent. De plus,

la quasi-colinéarité des métriques (coefficients ≥ 0.95 hors N_{tot}) justifie la sélection de $\mathcal{E}_{p,max}$ comme observable pilier, réduisant la dimension de sortie pour le contrôle optimal.

En conclusion, ces résultats démontrent que, dans nos conditions expérimentales, le contrôle spatial du front d'onde est beaucoup plus influent que le modelage temporel en termes d'impact sur les performances du TNSA, et qu'$\mathcal{E}_{p,max}$ peut être utilisé comme indicateur robuste et unique pour piloter des stratégies d'optimisation basées sur le ML. Cela ouvre la voie à un contrôle en temps réel, à 0.625 Hz, où un algorithme d'optimisation séquentielle ajuste un sous-ensemble d'actionneurs (centré sur l'actionneur 0) tout en surveillant $\mathcal{E}_{p,max}$ comme critère unique, assurant une convergence rapide et une robustesse accrue. À plus long terme, l'approche est généralisable au jeu complet des 48 actionneurs et à d'autres géométries/épaisseurs de cibles, avec des perspectives d'augmentation de stabilité, de rendement énergétique et de répétabilité des sources laser-plasma.

Chapitre 5 — Optimisation bayésienne assistée par modèles de substitution pour l'accélération laser de protons

Ce chapitre présente une stratégie d'optimisation multi-étapes (*Multi-Step Bayesian Optimization*, MSBO) destinée à améliorer les performances des faisceaux de protons générés dans le régime TNSA par façonnage adaptatif du front d'onde au moyen d'un MD (Catrix et al., 2025). L'approche exploite un modèle d'apprentissage automatique basé sur des forêts aléatoires pour prédire la relation entre tensions appliquées aux actionneurs du MD et métriques du faisceau de protons, suivi d'une OB pour identifier une des configurations optimales dans un espace de paramètres à 20 dimensions (actionneurs 0 à 19) — une réduction de dimension guidée par des considérations d'efficacité computationnelle et par la zone d'impact des actionneurs les plus centraux sur le faisceau laser.

L'installation expérimentale (Figure 5.1) repose sur le laser Ti :saphir 150 TW de l'*ALLS*, délivrant 3.2 J en 22 fs, focalisé par une parabole hors-axe $f/3$ sur des cibles d'aluminium de 4.5 µm. Le contraste temporel du faisceau ($< 10^{-10}$ à -100 ps, $< 10^{-6}$ à -3 ps) assure des conditions d'interaction optimales et reproductibles. Le diagnostic repose sur un détecteur de type parabole de Thomson couplé à un MCP, permettant la mesure en temps réel du spectre de protons tir par tir. L'ensemble est complété par un MD bimorphe à 48 actionneurs (Figure 5.3), couplé à un analyseur de front d'onde *SID4* pour caractériser le front d'onde du laser, et à une caméra CCD pour visualiser la tache focale. La combinaison d'un porte-cibles multi-échantillons (jusqu'à 400 cibles) et de la stabilité intrinsèque de la chaîne laser (fluctuations de 2.5% RMS) garantit la possibilité de collecter des jeux de données statistiques fiables, condition essentielle à l'application de méthodes d'optimisation basées sur l'apprentissage automatique dans un contexte de cadence expérimentale limitée (0.625 Hz).

La méthode MSBO se décompose en quatre étapes (Figures 5.4 – 5.11). Après une phase d'échantillonnage initial par *Latin Hypercube Sampling* (Helton et al., 2003) des 20 premiers actionneurs (Figure 5.5), un modèle RF est entraîné sur les données expérimentales et comparé à d'autres modèles (linéaire, perceptron multicouche, *k-nearest neighbors*, *Support Vector Regression*). Le modèle RF se distingue avec une précision élevée (R^2 = 0.97, RMSE = 0.24 MeV), confirmant sa capacité à capturer les relations non linéaires complexes et bruitées inhérentes aux interactions laser–plasma. L'analyse des importances relatives (Figure 5.7) révèle l'influence dominante de l'Actionneur 0 (défocus global), avec une corrélation de 0.86 avec $\mathcal{E}_{p,max}$, alors

que les autres actionneurs présentent des corrélations faibles et peu interdépendantes. L'OB guidée par ce modèle identifie rapidement des configurations efficaces, atteignant 90 % des performances optimales en $\sim$30 tirs et convergeant en 100–150 tirs (Figure 5.18), soit un gain considérable par rapport aux méthodes classiques.

Les résultats expérimentaux (Figure 5.12) montrent une augmentation de $\mathcal{E}_{p,max}$ de 72 % par rapport à la configuration du MD où tous les actionneurs sont à 0 V, et de 24 % par rapport à la configuration du MD donnée par l'analyseur de front d'onde qui minimise les aberrations pour converger vers un front d'onde plan, atteignant (6.64 ± 0.30) MeV. Fait marquant, la configuration optimale n'est pas celle correspondant à un front d'onde plan (Figures 5.13.A–D) et n'atteint pas l'intensité focale maximale, ce qui suggère que l'optimisation agit sur les conditions de couplage laser–plasma plutôt que sur la seule concentration d'énergie. En d'autres termes, ce travail démontre que maximiser la qualité optique classique d'un faisceau laser, en maximisant par exemple le *Strehl ratio*, n'implique pas nécessairement une maximisation de l'efficacité du mécanisme TNSA, un résultat qui bouscule l'approche traditionnelle de l'optique adaptative appliquée aux lasers de puissance.

Une analyse modale basée sur la décomposition de Laguerre–Gauss (Figures 5.15–5.17) a révélé une structure radiale complexe conservant une forte symétrie cylindrique ($|\ell| = 0$ dominant à 91%), avec 95% de l'énergie concentrée dans les 10 premiers modes. Cette compacité modale, malgré le caractère principalement non-gaussien du faisceau, suggère que la MSBO génère un profil structuré en intensité, qui optimise le transfert d'énergie vers les électrons chauds tout en minimisant les pertes latérales. Ce résultat rejoint les travaux récents sur les faisceaux porteurs de moment angulaire orbital, où ce caractère joue un rôle dans la collimation des protons et la limitation des pertes transverses (Willim et al., 2023).

L'étude de la phase spatiale (front d'onde) par décomposition en modes de Zernike (Figures 5.19–5.21, Tableau 5.4) montre que la configuration MSBO exploite volontairement des aberrations de bas et moyen ordre (coma, tilt, astigmatisme), contrairement à la correction donnée par l'analyseur de front d'onde qui les minimise. Malgré un *Strehl ratio* réduit (0.179 contre 0.692), la configuration MSBO conduit à une meilleure efficacité de conversion laser–protons (0.74% contre 0.61%). Ce constat illustre un paradigme intéressant : certaines aberrations optiques, loin d'être purement nuisibles, deviennent des leviers de contrôle pour améliorer l'absorption laser et la génération d'électrons chauds.

Ces observations, confrontées aux lois d'échelle du TNSA (Fuchs et al., 2006; Robson et al., 2007; Flippo et al., 2008; Zeil et al., 2010), confirment que l'amélioration obtenue ne provient pas d'une intensité focale accrue, mais bien d'un couplage laser–plasma optimisé. Dans notre configuration MSBO, la température des électrons chauds atteint $T_{hot} = 1.58$ MeV, soit une valeur supérieure à celle prédite par les seules lois d'échelle en intensité, et nettement supérieure à celle mesurée avec un faisceau corrigé (1.16 MeV). Cela montre que le front d'onde optimal pour le TNSA n'est pas nécessairement celui qui maximise l'intensité, mais celui qui maximise l'efficacité d'absorption et la symétrie du champ de gaine.

En définitive, la MSBO démontre la pertinence des approches de façonnage non conventionnelles du front d'onde pour repousser les performances des sources de protons laser-plasma. L'approche est particulièrement adaptée aux installations où la mesure du front d'onde à pleine puissance est impossible, et où le nombre de tirs disponibles est limité : elle offre un cadre général pour des

optimisations rapides, sûres et robustes. À plus long terme, la méthodologie pourrait être étendue à un espace de paramètres élargi incluant plus d'actionneurs, l'énergie laser, la polarisation ou le type de cible ouvrant la voie à des stratégies d'optimisation multidimensionnelles pour l'ensemble des régimes d'interaction laser-plasma.

Conclusion

Ce travail démontre que l'optimisation guidée par ML n'est pas seulement un concept théorique, mais bien une stratégie pratique et efficace pour améliorer les performances de l'accélération de protons par TNSA dans des conditions expérimentales réalistes. Au-delà des gains spécifiques obtenus sur l'énergie maximale des protons, le cadre méthodologique modulaire développé ici constitue une voie évolutive vers un contrôle adaptatif, fondé sur les données, des interactions laser–plasma à haute intensité. L'extension de cette approche vers un fonctionnement en temps réel, combinée à l'intégration d'algorithmes d'optimisation complémentaires et à son application à d'autres régimes d'interaction, constituerait un levier déterminant pour accélérer la transition des expériences de démonstration de principe vers des sources laser–plasma autonomes, stables et pleinement opérationnelles, adaptées aux exigences concrètes des applications scientifiques, industrielles et sociétales.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

AI	Artificial Intelligence
ALLS	Advanced Laser Light Source
ANN	Artificial Neural Network
BO	Bayesian Optimization
BOA	Break-Out Afterburner Acceleration
CEP	Carrier Envelope Phase
CNN	Convolutional Neural Network
CPA	Chirped Pulse Amplification
DL	Deep Learning
DM	Deformable Mirror
EI	Expected Improvement
FWHM	Full Width at Half Maximum
GA	Genetic Algorithm
GDD	Group Delay Dispersion
GP	Gaussian Process
GPR	Gaussian Process Regression
HEDP	High Energy Density Physics
HPC	High-Performance Computing
HPL	High Power Laser
HPO	Hyperparameter Optimization
HRR	High Repetition Rate
INRS	Institut national de la recherche scientifique
LDPA	Laser-Driven Particle Acceleration
LPI	Laser-Plasma Interaction
LPWA	Laser-Plasma Wakefield Acceleration
LSTM	Long Short-Term Memory
ML	Machine Learning
NN	Neural Network
PET	Positron Emission Tomography
PI	Probability of Improvement

PIC	Particle-In-Cell
PSO	Particle Swarm Optimization
PtV	Peak-to-Valley
RF	Random Forest
RL	Reinforcement Learning
RMS	Root Mean Square
RNN	Recurrent Neural Network
RPA	Radiation Pressure Acceleration
SA	Simulated Annealing
SMA	Shockwave Mediated Acceleration
SVR	Support Vector Regression
TNSA	Target Normal Sheath Acceleration
TOD	Third Order Dispersion
TPOT	Tree-based Pipeline Optimization Tool
UCB	Upper Confidence Bound

INTRODUCTION

The advent of high-power lasers, which can produce compact bunches of coherent photons, has opened up possibilities of laser-based particle acceleration. The original idea of laser-driven particle acceleration was initially brought up by Toshiki Tajima and John M. Dawson in 1979 (Tajima et al., 1979), at a timely moment with regards to the advent of the Chirped Pulse Amplification (CPA) technique, which would occur six years after (Strickland et al., 1985) developed by the 2018 Nobel laureates Donna Strickland and Gérard Mourou. These two events further triggered the development of laser-based electron and ion acceleration. The acceleration of electrons has been studied since the 70s-80s through wakefields and has very recently made a breakthrough by reaching the GeV limit (Leemans et al., 2006) in mono-energetic bunches (Geddes et al., 2004; Faure et al., 2004; Mangles et al., 2004). The first evidence of proton production by laser-target interaction was observed in experiments between 1978 and 1982, yielding very low maximum proton energies of $\sim$0.56 MeV. Among the various mechanisms explored since Toshiki Tajima and John Dawson first proposed laser-driven particle acceleration in their seminal paper (Tajima et al., 1979), the Target Normal Sheath Acceleration (TNSA) regime has emerged as one of the most widely implemented and experimentally accessible approaches for ion acceleration. The field was launched by three landmark experiments in 2000, which reported the generation of multi-MeV proton beams from solid targets irradiated by high-intensity lasers (Snavely et al., 2000; Clark et al., 2000; Maksimchuk et al., 2000). These findings demonstrated the potential of compact, tabletop laser systems to rival large-scale particle accelerators.

At the heart of this field lies the interaction of ultra-intense laser pulses, typically exceeding 10^{18} W.cm^{-2}, with solid, gaseous, or pre-ionized plasma targets. These interactions produce extreme physical conditions, including relativistic electrons, highly non-equilibrium plasmas, and strong electromagnetic fields. Rapid ionization of the target leads to charge separation and the formation of space-charge fields capable of sustaining gradients on the order of tens of TV/m over micrometer-scale distances (Romagnani et al., 2005). These gradients can accelerate ions, especially protons, to tens of MeV over just a few microns, far shorter than the meter-scale acceleration lengths required by conventional accelerators.

Ion acceleration mechanisms, including TNSA (Wilks et al., 2001), Radiation Pressure Acceleration (RPA) (Robinson et al., 2008), Collisionless Shock Acceleration (CSA) (Silva et al., 2004; Haberberger et al., 2012), Magnetic Vortex Acceleration (MVA) (Nakamura et al., 2010; Park et al., 2019), and Break-Out Afterburner (BOA) (Yin et al., 2006, 2007), rely on different physical processes. The

efficiency and spectral characteristics of each regime depend sensitively on laser parameters such as pulse duration, intensity, contrast, wavefront quality, focal geometry, and polarization state.

The development of TW- and PW-class laser systems has expanded the frontier of compact and tunable ion accelerators. These sources offer substantial advantages over conventional radiofrequency-based systems, including ultrashort beam durations, micron-scale source sizes, high peak currents, and excellent laminarity. Facilities such as *ELI-Beamlines* (ELI - Extreme Light Infrastructure, 2011) and *Apollon* (Papadopoulos et al., 2016) (*Laboratoire pour l'utilisation des lasers intenses*, France) are enabling higher ion energies, better beam stability, and increased repetition rates. Nonetheless, a critical hurdle for practical applications is the increase in laser repetition rate while preserving ultra-high intensities. Recent advances in thin-disk laser technology (Giesen et al., 2007; Baer et al., 2010; Saraceno et al., 2019) and fiber-based architectures (Richardson et al., 2010; Zervas et al., 2014) are paving the way toward compact, high-average-power, and reliable laser systems for next-generation ion sources.

In parallel, Machine Learning (ML) and data-driven optimization techniques, such as Bayesian Optimization (BO), have emerged as powerful tools to enhance the performance and efficiency of laser-plasma acceleration systems (Gonoskov et al., 2019; Lin et al., 2021; Döpp et al., 2023; Roussel et al., 2024). Given the high dimensionality of the parameter space — including laser parameters, target configurations, and diagnostic settings — and the high cost of experimental shots, ML offers an efficient strategy for intelligent exploration and optimization with minimal iterations. Surrogate models, such as Gaussian Processes (GP), Random Forests (RF), or neural networks (NNs), can be trained on experimental or simulated datasets to accurately predict key particle beam characteristics, including maximum energy, total yield, and beam divergence, enabling real-time decision-making and closed-loop optimization (Ferran Pousa et al., 2023). Recent studies have demonstrated that incorporating ML into laser-driven ion acceleration workflows can significantly reduce experimental time, improve reproducibility, and uncover non-intuitive control strategies that would be difficult to identify using traditional approaches (Shalloo et al., 2020; Djordjević et al., 2021b; Loughran et al., 2023; Beier et al., 2024; McQueen et al., 2025). As laser systems grow more complex and tunable, the integration of ML will likely play a central role in enabling autonomous and adaptive control of next-generation laser-driven ion sources.

Worldwide, an increasing number of laboratories are commissioning laser-driven proton beamlines (Giulietti et al., 2007; Yogo et al., 2009; Agosteo et al., 2014; Cirrone et al., 2015; Kim et al., 2016; Wagner et al., 2016; Higginson et al., 2018; Obst-Huebl et al., 2018; Barberio et al., 2020), supporting a rapidly expanding range of applications currently under active investigation, including :

— Medical applications : such as hadron therapy, radiobiology, and Flash therapy (Bulanov et al., 2002; Malka et al., 2004; Flacco et al., 2024),

— Material science : including ion-induced structural studies (Sciscio et al., 2017; Barberio et al., 2017a,b, 2018a; Puyuelo-Valdes et al., 2021),

— Astrophysics (Patel et al., 2003; Albertazzi et al., 2014)

— and many more.

The current doctoral thesis aims to enhance our understanding and control of laser-driven ion acceleration, with a particular focus on improving the reproducibility and performance of proton beams in the TNSA regime for unlocking bottlenecks of some applications. The work is carried out using the 150 TW Ti : Sapphire laser system at the *Advanced Laser Light Source* (*ALLS*) facility in Varennes, Canada. Key objectives include enhancing laser-driven proton beam reproducibility through the deployment of high-stability target delivery systems, determining optimal laser–target interaction conditions, and applying ML strategies to enable efficient experimental optimization. This research combines advanced diagnostics, adaptive optical control, and data-driven modeling to develop reliable, high-performance laser-plasma sources for compact accelerator applications. This integrated approach deepens our understanding of laser–matter interactions while enabling data-driven optimization of next-generation ion sources.

This doctoral thesis is structured as follows :

Chapter 1 introduces the fundamental physical principles and technological innovations underpinning laser-driven particle acceleration. It begins by tracing the development of high-power lasers and the breakthrough of CPA, which revolutionized ultrafast science. The chapter then explores the electromagnetic foundations necessary to understand the dynamics of laser-plasma interactions and ion acceleration. Building on this foundation, it presents the key mechanisms of laser-driven particle acceleration by intense lasers, covering both electron and ion acceleration. Particular emphasis is placed on TNSA, with a detailed discussion of hot electron generation, sheath field formation, and the resulting ion acceleration.

Chapter 2 explores the transformative role of ML in advancing laser-driven particle acceleration. It begins by defining key concepts in ML and tracing its recent emergence in plasma physics research. The chapter then delves into the inherent complexity of laser-plasma interactions, examining nonlinear optical phenomena, heating mechanisms, plasma instabilities, and the challenges of experimental modeling. A key focus is placed on the high-dimensional parameter space governing laser-plasma acceleration, including laser, target, pre-plasma, and environmental parameters. The chapter then reviews various ML techniques for optimization such as neural networks, support vector regression, random forests, Gaussian process regression, and reinforcement learning, highlighting their potential for improving beam quality and experimental efficiency. A discussion section evaluates the broader impact of ML on experimental workflows and outlines insights gained from data-driven approaches. Finally, future perspectives emphasize the integration of ML with real-time feedback control and high-performance computing for next-generation acceleration platforms.

Chapter 3 demonstrates the capability of the *ALLS* 150 TW ion acceleration beamline to reliably generate stable proton beams over hundreds of consecutive shots in the TNSA regime, using a custom-designed 400-target holder accommodating various foil materials and thicknesses. Despite some shot-to-shot variability, the proton yield ($<$26%) and proton maximum energy ($<$7%) show low shot-to-shot fluctuations compared to other facilities. These results are particularly significant

in the 1-5 MeV range for different target types. This setup ensures robust and reproducible proton beam generation, offering improved stability over previous similar studies and facilitating rapid target alignment, thus providing a strong foundation for further optimization and development of laser-plasma ion sources.

Chapter 4 presents a comprehensive study of proton acceleration in the TNSA regime, focusing on spatio-temporal laser beam shaping. 20 out of 48 actuators of a bimorph deformable mirror (DM) and an acousto-optic programmable dispersive filter were used to independently control the *ALLS* laser's spatial wavefront and spectral phase, respectively. By evaluating the response of key laser-driven proton beam metrics to these adjustments, we demonstrate that spatial shaping plays a predominant role in enhancing laser-driven proton beam quality, while temporal phase modulation has a comparatively minor effect when using 4.5-µm-thick aluminum targets. In this context, the maximum proton energy emerges as a particularly relevant optimization metric, as it strongly correlates with other beam characteristics and thus serves as a reliable proxy for the overall TNSA efficiency.

Finally, **Chapter 5** introduces an innovative ML-based strategy to enhance the performance of laser-driven proton acceleration, with a focus on maximizing proton energy in the TNSA regime. By optimizing 20 out of 48 actuators of a bimorph DM using a multi-step BO (MSBO) framework guided by a random forest surrogate model, a 72% increase in the maximum proton energy was achieved. Starting from zeroed DM actuator voltages, our method identified an optimal configuration using 20 out of 48 actuators, requiring fewer than 150 experimental laser shots. Our method surpassed conventional wavefront correction by 24%, which typically minimizes aberrations to converge toward a flat wavefront by leveraging real-time feedback from a wavefront sensor. This data-driven approach, which integrates advanced wavefront control, challenges the conventional emphasis on aberration correction aimed at producing a flatter wavefront in laser-driven ion acceleration. This work demonstrates that surrogate-assisted, data-guided optimization can significantly accelerate the tuning of high-energy laser–plasma interactions, even when diagnostic access is limited. By rapidly converging toward optimal conditions in complex parameter spaces, the method opens a pathway to widespread implementation of ML-based control at experimental facilities and to more stable, higher-performance laser-driven ion sources.

1 HIGH-POWER LASER TECHNOLOGIES AND ELECTROMAGNETISM AS FOUNDATIONS OF PARTICLE ACCELERATION

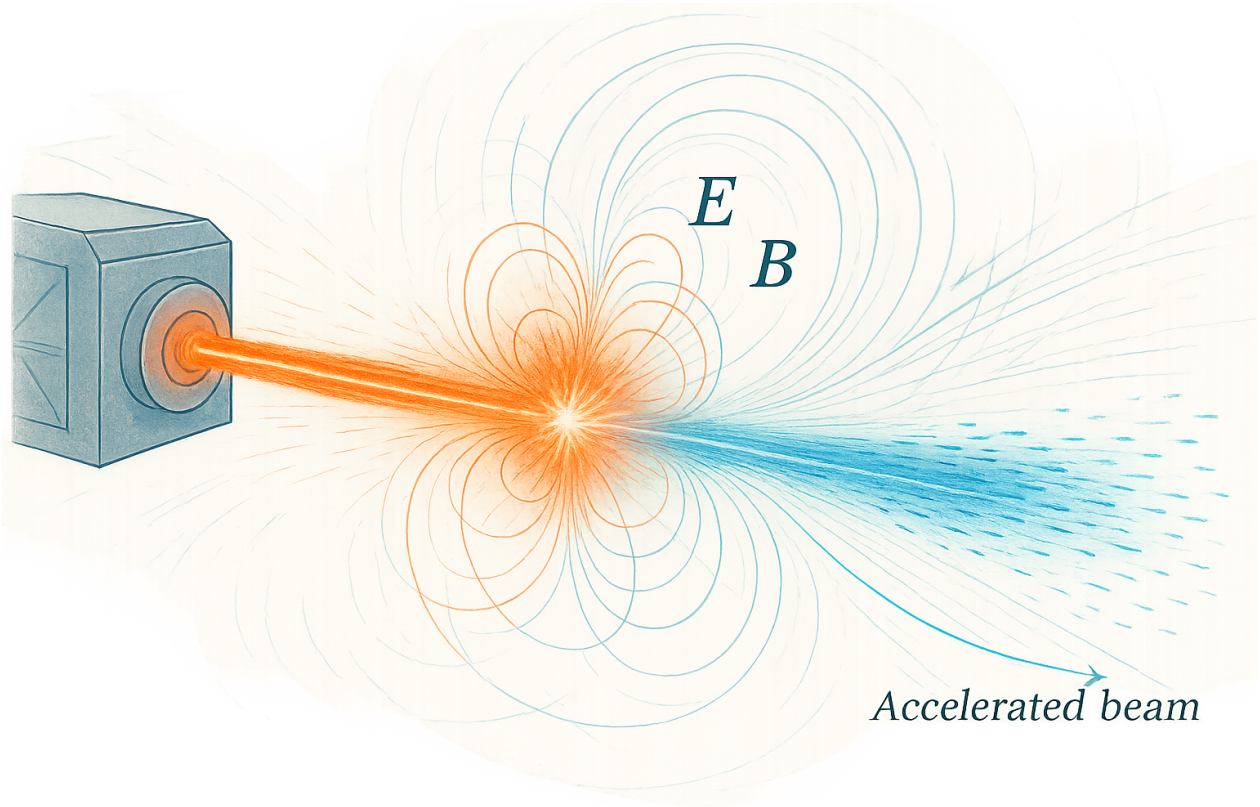

Intense laser–matter interaction generates strong electromagnetic fields that accelerate ions into a forward-directed high-energy beam. Original image created by the author (Elias Catrix). No external permission is required; use authorized by the author.

Foreword — This chapter serves as a comprehensive foundation for understanding the physics of laser-driven particle acceleration, with a particular emphasis on Target Normal Sheath Acceleration (TNSA). The content is structured to guide the reader from the historical and technological context of high-power lasers to the detailed electromagnetic and plasma-physics principles underlying ion acceleration.

We first introduce the historical development of Chirped Pulsed Amplification and its role in propelling high-intensity laser science, before revisiting the electromagnetic framework that governs plasma interactions. We then present the main acceleration mechanisms for both electrons and ions, starting with a concise overview of Laser Wakefield Acceleration for electrons, followed by a detailed discussion of TNSA. Complementary mechanisms such as Radiation Pressure Acceleration, Collisionless Shock Acceleration, Magnetic Vortex Acceleration, and Break-Out Afterburner are also introduced to provide a broader context of high-intensity ion acceleration. Finally, we examine the practical aspects of laser–plasma experiments, including target engineering, beam diagnostics, and the scaling laws that guide experimental optimization.

Together, these elements provide a conceptual and technical foundation for the experimental strategies and optimization methods explored in the chapters that follow.

1.1 High-power lasers and Chirped Pulse Amplification : Revolutionizing ultrafast science

The development of high-power laser systems has fundamentally transformed the study of laser-matter interactions, enabling groundbreaking advances in plasma physics and particle acceleration. At the heart of this transformation lies the technique of Chirped Pulse Amplification (CPA). In their seminal 1985 paper (Strickland et al., 1985), Donna Strickland and Gérard Mourou proposed a method to circumvent the limitations of traditional laser amplification that led to nonlinear distortions or damage due to peak intensities. The CPA technique consists of three essential steps as shown in Figure 1.1 :

1. **Stretching** the ultrashort pulse temporally (typically from femtoseconds to picoseconds) using a dispersive element (e.g., a pair of gratings), reducing its peak power.
2. **Amplifying** the stretched pulse using a gain medium (e.g., Ti : Sapphire).
3. **Recompressing** the pulse back to its original (or nearly original) duration, recovering the high peak intensity.

The CPA technique allows the amplification of pulses to joule-level energies while maintaining femtosecond durations, leading to peak powers to the TW level (10^{12} W) and up to 10 PW (10×10^{15} W). The most powerful lasers currently operational are a pair of 10 PW lasers located in the *Extreme Light Infrastructure* (*ELI*) - Nuclear Physics facility in Măgurele, Romania (Radier et al., 2022). In the UK, a laser called *Vulcan 20-20* will become the most powerful (20 PW) laser in the

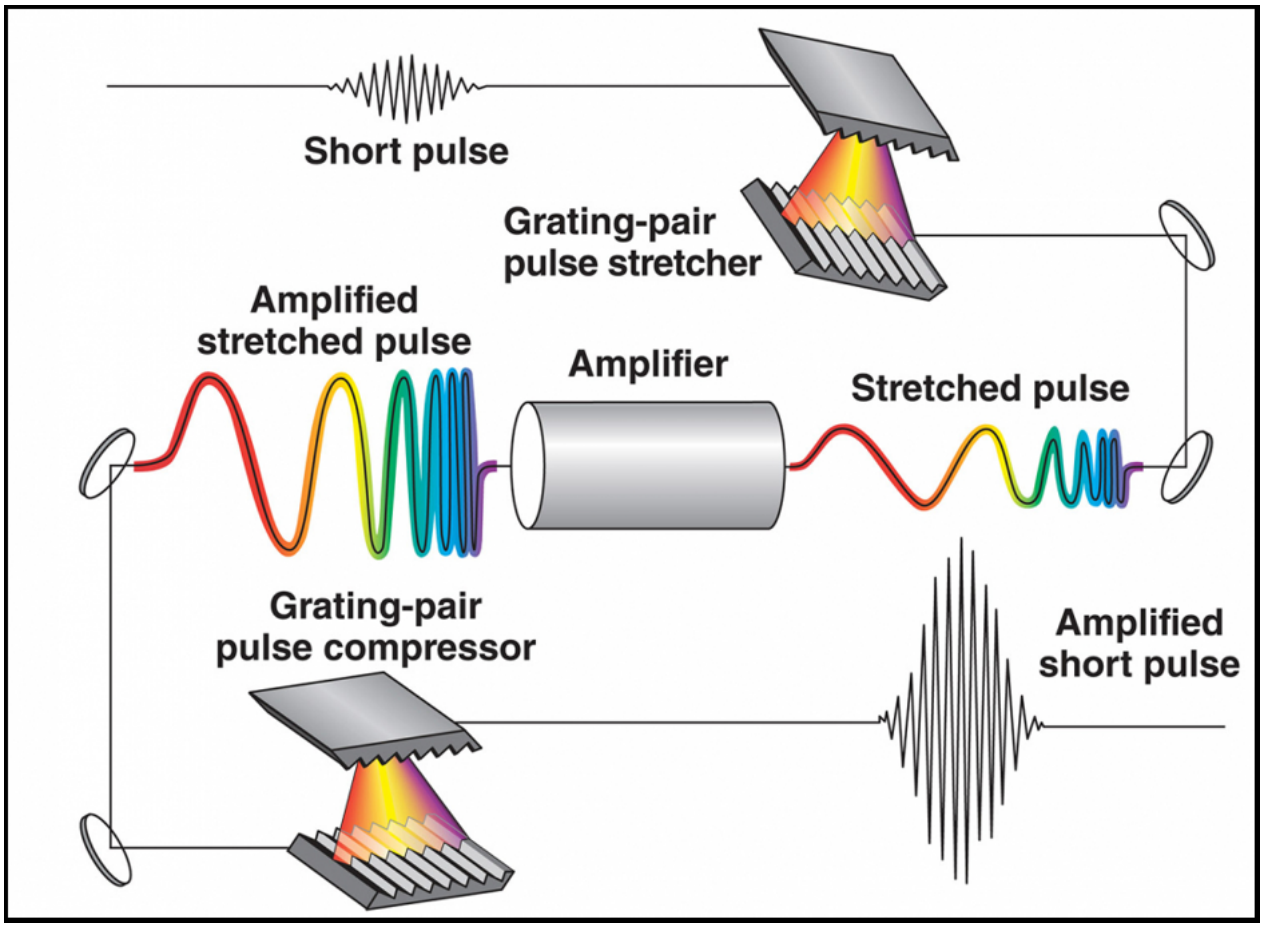

FIGURE 1.1 : Principle of Chirped Pulse Amplification.

This figure has been adapted from Strickland et al. (1985), with the publisher's permission.

world when it is finished in 2029 (Baraniuk, 2023). Figure 1.2 provides a global overview of the major high-intensity laser labs in operation or development as of 2020.

Building on these international efforts, the development of PW-class systems has a rich history, marked by early milestones that paved the way for today's most powerful laser infrastructures. In 1994, Barty et al. (1994) demonstrated the feasibility of extending CPA systems to PW-class peak powers (e.g., $\sim$15 J in $\sim$15 fs) using available technologies. The Nova Petawatt Laser, for example, reached 1 PW in 1996 at the *Lawrence Livermore National Laboratory* (*LLNL*) (Barty et al., 2004). This milestone led to a rapid global expansion of ultrafast, high-intensity laser systems, many of which now serve as platforms for fundamental research in laser-plasma interactions, attosecond science, and particle acceleration. Ti : Sapphire CPA lasers dominate due to their broad gain bandwidth, enabling sub-30 fs pulse durations (Joyce et al., 2010; Dabu, 2018). This enables exceptional temporal resolution, making them ideal for investigating ultrafast phenomena. Other systems based on Yb :glass (Taira et al., 1997; Südmeyer et al., 2009; Morrissey et al., 2017; Dietz et al., 2020), or Optical Parametric CPA, commonly named OPCPA (Ross et al., 2002; Xu

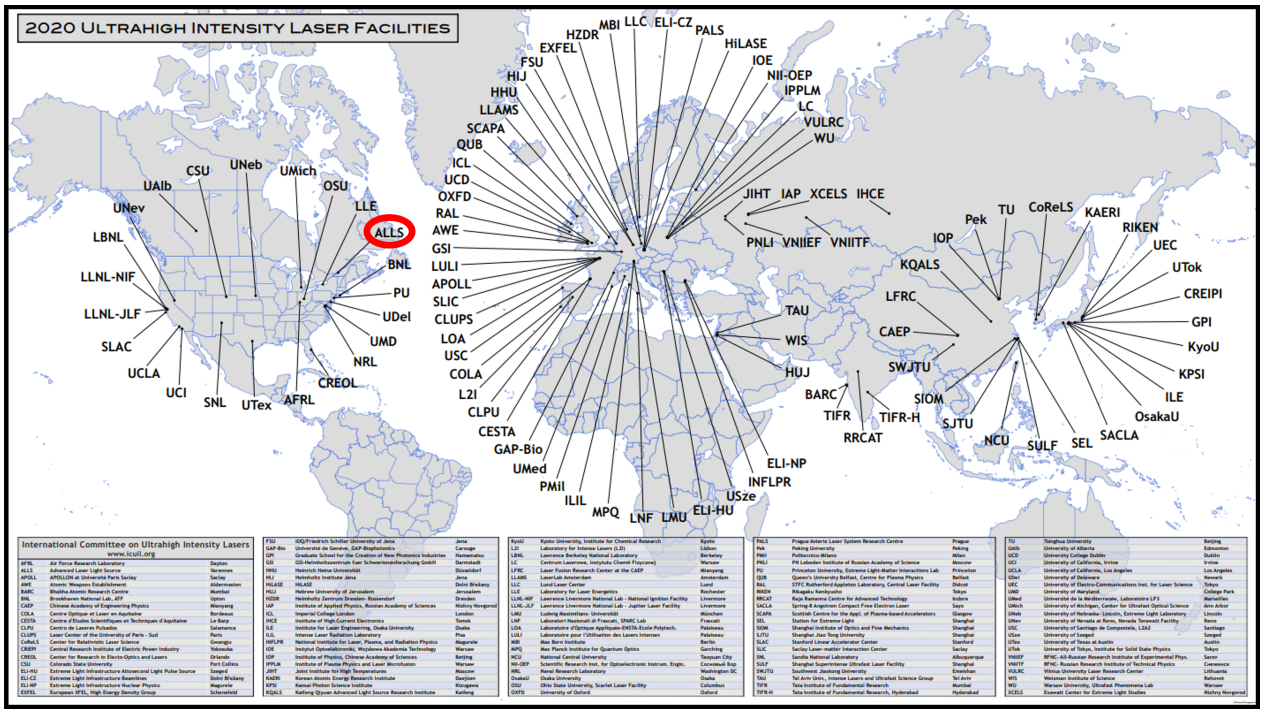

FIGURE 1.2 : Worldwide map of ultrahigh-intensity lasers

Worldwide distribution of ultrahigh-intensity laser facilities as of 2020. This figure has been adapted from International Committee on Ultrahigh Intensity Lasers (2020), with the publisher's permission.

et al., 2013; Zeng et al., 2017; Xiao et al., 2021; Wang et al., 2022) have emerged for different pulse formats and repetition rates. Figure 1.3 demonstrates the chart of the historical development of laser systems with high pulse power and intensity. Although facilities like *ELI*, *IZEST*, and C^3 are among the most advanced high-intensity laser projects, they have not yet reached the 10^{25} W.cm^{-2} threshold due to several limiting factors. While their energy and compression capabilities suggest that this intensity is within reach, practical constraints, such as optical damage thresholds, imperfect vacuum conditions, and focusing limitations, prevent its full realization. At these intensities, focusing PW-class pulses to micrometer-sized spots without introducing aberrations or damaging optics remains a major engineering challenge. Thus, despite theoretical potential, reaching 10^{25} W.cm^{-2} still requires substantial advances in laser technology, targetry, and vacuum infrastructure.

One of the central challenges in scaling CPA systems to higher pulse energies and peak powers lies in mitigating various nonlinear and thermal effects that arise during amplification. Among these, gain narrowing can significantly limit the spectral bandwidth, thus lengthening the recompressed pulse and reducing temporal resolution (Perry et al., 1994). Additionally, high average powers can induce thermal lensing and stress in optical components, threatening system stability and beam quality (Brida et al., 2007). To address these limitations, several technical innovations have been developed. OPCPA offers an alternative to traditional gain media by utilizing nonlinear crystals for parametric gain, enabling amplification over broader spectral bandwidths and with superior temporal contrast (Dubietis et al., 1992). Additional techniques, such as double CPA architectures

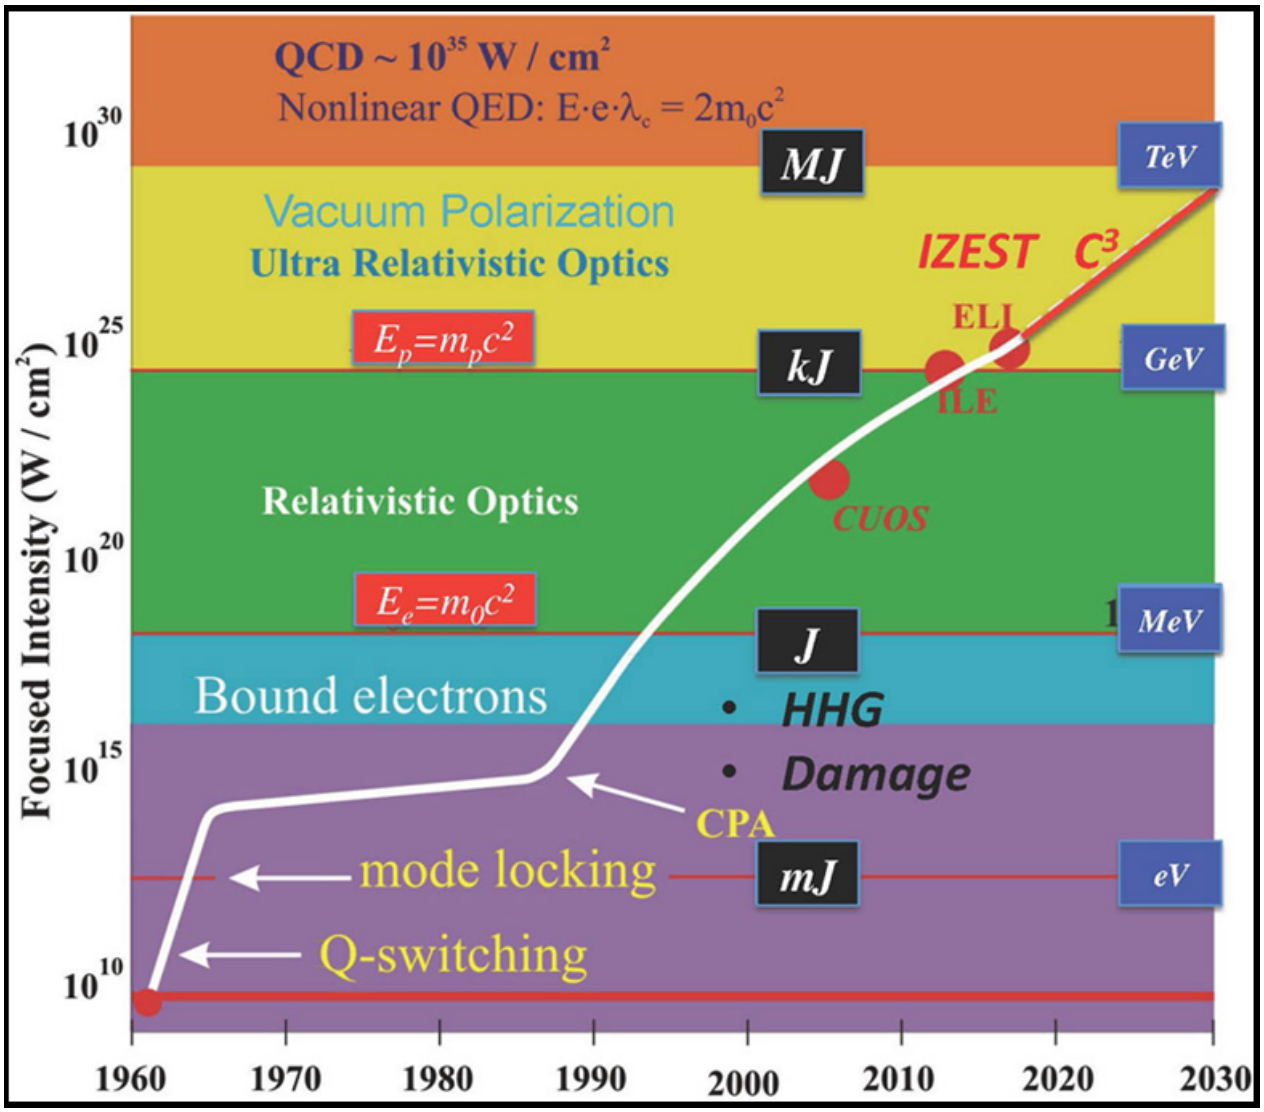

FIGURE 1.3 : Historical evolution of high-power laser systems.

Timeline showing the historical evolution and projected development of pulsed high-power laser systems as of 2012. This figure has been adapted from Mourou et al. (2012), with the publisher's permission.

and cross-polarized wave (XPW) generation, as implemented at the *Advanced Laser Light Source* (*ALLS*) laser facility, have been employed to enhance the temporal contrast of ultrashort laser pulses significantly. These methods are particularly effective in suppressing pre-pulses and pedestal energy, which is crucial for ensuring high-contrast conditions in laser-plasma interaction experiments (Jullien et al., 2005; Levy et al., 2008; Fourmaux et al., 2011b). Moreover, thin-disk amplifiers (Giesen et al., 2007; Baer et al., 2010; Saraceno et al., 2011; Jung et al., 2016; Saraceno et al., 2019) and fiber-based (Richardson et al., 2010; Zervas et al., 2014) CPA configurations provide efficient thermal management and are well-suited for high-repetition-rate operation.

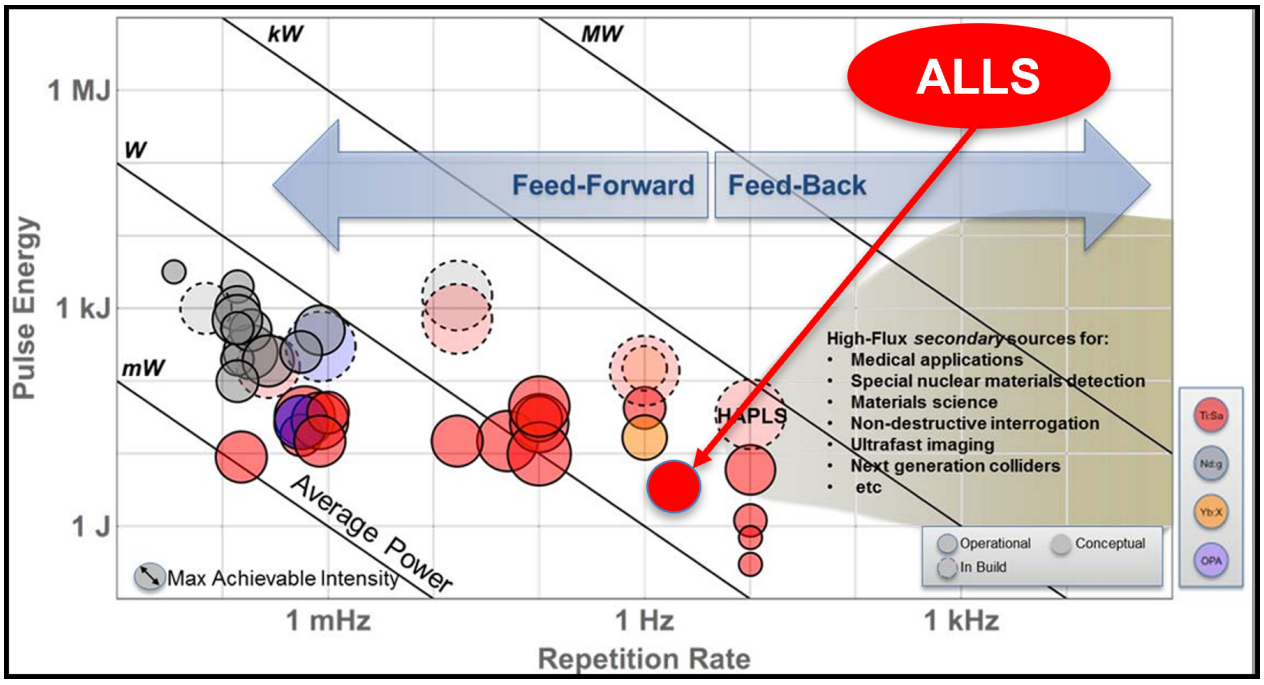

FIGURE 1.4 : Pulse energy versus repetition rate of high-power lasers.

This figure has been adapted from Siders et al. (2016), with the authors' permission.

The field has progressively evolved toward higher repetition rates and increased average powers (Figure 1.4), as exemplified by systems such as *LLNL*'s High-repetition-rate Advanced Petawatt Laser System (HAPLS). This development marks a pivotal transition toward feedback-optimized, high-performance laser platforms designed to serve both fundamental research and societal applications. These advancements collectively contribute to the ongoing evolution of CPA systems, pushing the frontiers of ultrafast, high-intensity laser science. One of the most prominent applications is in laser–plasma acceleration, where intense femtosecond laser pulses are focused into an underdense plasma to generate large electric fields capable of accelerating electrons to GeV energies over centimeter-scale distances, a process known as Laser Wakefield Acceleration (LWFA). In contrast, Target Normal Sheath Acceleration (TNSA) occurs when high-intensity laser pulses interact with solid targets and form overdense plasmas, producing ultra-short, high-energy bursts of protons and heavier ions, with applications ranging from inertial confinement fusion to ultrafast radiography. In the realm of attosecond science, CPA lasers are fundamental in driving high-harmonic generation, enabling the generation of coherent extreme ultraviolet and soft X-ray pulses with attosecond durations (Krausz et al., 2009). These tools allow for real-time probing of electron dynamics in atoms, molecules, and condensed matter. Furthermore, CPA-based laser systems are increasingly used to investigate warm dense matter, a regime relevant to planetary physics and inertial fusion, where matter exists at high energy densities but near-solid densities (Glenzer et al., 2009). The availability of PW-class beamlines at major facilities such as *ELI* (Danson et al., 2019), the *LaserNetUS* network (LaserNetUS, 2025) or *Laserlab Europe* (Laserlab-Europe Consortium, 2025)

has democratized access to these cutting-edge experimental platforms, enabling broad international collaboration and rapid scientific progress across disciplines.

1.2 Electromagnetic foundations of laser-driven particle acceleration

Electromagnetic fields are at the heart of laser-plasma accelerators. They accelerate, focus, and guide charged particles, and they are responsible for the stability as well as the instability of particle beams. Their range extends from constant fields up to very fast-changing fields with frequencies of many GHz.

Long ago, electricity and magnetism were well-known separate phenomena. The birth of electromagnetism began with the discovery of Hans Christian Ørsted (1820), who found that an electric current is always associated with a magnetic field. Later, Michael Faraday (1831) discovered electromagnetic induction, which led to the creation of electric fields by changing magnetic fields. Electromagnetism was born. James Clerk Maxwell (1864) extended and completed this work with four equations that relate the electric field $\vec{E}$ and the magnetic field $\vec{H}$, together with the electromagnetic Lorentz force. The four equations are :

Equation 1.1 : Ampère–Maxwell law

$$\oint \vec{H}(\vec{r},t)\cdot d\vec{s} = \iint \vec{J}(\vec{r},t)\cdot d\vec{A} + \frac{d}{dt}\iint \vec{D}(\vec{r},t)\cdot d\vec{A} \tag{1.1}$$

Equation 1.2 : Faraday's law of induction

$$\oint \vec{E}(\vec{r},t)\cdot d\vec{s} = -\frac{d}{dt}\iint \vec{B}(\vec{r},t)\cdot d\vec{A} \tag{1.2}$$

Equation 1.3 : Gauss's law

$$\iint \vec{D}(\vec{r},t)\cdot d\vec{A} = \iiint \rho(\vec{r},t)\,dV \tag{1.3}$$

Equation 1.4 : Gauss's law for magnetism

$$\iint \vec{B}(\vec{r},t)\cdot d\vec{A} = 0 \tag{1.4}$$

with :

— $\vec{E},\vec{H}$ the electric and magnetic fields,

— $\vec{D},\vec{B}$ the electric displacement and the magnetic induction, which are responsible for the effects of the material on the fields,

— $\vec{J}$ the electric current density,

— ρ the electric charge density.

The term $\iint \vec{J}(\vec{r},t) \cdot d\vec{A}$ includes all currents going through the area $\vec{A}$, which may consist of :

Equation 1.5 : Conduction current density

$$\vec{J_c}(\vec{r},t) = \kappa \vec{E}(\vec{r},t) \tag{1.5}$$

Equation 1.6 : Convection current density

$$\vec{J_{cv}}(\vec{r},t) = \rho(\vec{r},t)\,\vec{v}(\vec{r},t) \tag{1.6}$$

Equation 1.7 : Additional (induced or external) current density

$$\vec{J_i}(\vec{r},t) \tag{1.7}$$

where :

— $\vec{J_c}(\vec{r},t)$ is the conduction current density,

— $\vec{J_{cv}}(\vec{r},t)$ is the convection current density,

— $\vec{J_i}(\vec{r},t)$ represents additional current density (e.g., induced or external),

— κ is the electrical conductivity of the medium,

— $\vec{v}(\vec{r},t)$ is the velocity of charge carriers,

— $\vec{r}$ is the position vector,

— t is time.

The term $\iiint \rho(\vec{r},t)dV$ includes all charges in the volume V. Current and charge may have different distributions, e.g., point-like, lines, surfaces, and volumes.

In most cases, Maxwell's equations are used in their differential form, for which an extensive set of tools is available. They can be derived from equations 2.1 to 2.4 using Stokes' and Gauss' theorems :

Equation 1.8 : Ampère–Maxwell law (differential form)

$$\nabla \times \vec{H} = \vec{J} + \frac{\partial \vec{D}}{\partial t} \tag{1.8}$$

Equation 1.9 : Faraday's law of induction (differential form)

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \tag{1.9}$$

Equation 1.10 : Gauss's law (differential form)

$$\nabla \cdot \vec{D} = \rho \tag{1.10}$$

Equation 1.11 : Gauss's law for magnetism (differential form)

$$\nabla \cdot \vec{B} = 0 \tag{1.11}$$

It should be noted that Maxwell's theory is a continuum theory, in which all electromagnetic quantities must remain continuous. Apparent discontinuities, such as those found at material interfaces or in rapidly evolving plasma boundaries, are handled through the integral form of Maxwell's equations. In the context of laser-driven particle acceleration, whether of electrons or ions, both integral and differential formulations are used depending on the spatial and temporal scales of interest. However, due to the ultrafast, nonlinear, and highly localized nature of laser–plasma interactions, the differential form of Maxwell's equations is often more appropriate for theoretical modeling and numerical simulations. This is particularly relevant when resolving the dynamics of charge separation, sheath field formation, and field-driven acceleration, which are central to both electron and ion acceleration processes.

1.3 Fundamentals of laser-driven particle acceleration

This section outlines the fundamental physical principles underlying the main laser-driven particle acceleration mechanisms. Both electron and ion acceleration processes are governed by the interplay between ultra-intense laser fields and rapidly evolving plasma structures, where collective electromagnetic effects and charge separation dynamics play a central role.

A concise overview of laser-driven electron acceleration is first presented here. The objective is not to present an exhaustive treatment of the subject, as it lies outside the primary focus of this doctoral work, but rather to introduce the key concepts and major historical developments. For a more detailed discussion, the reader is referred to several excellent reviews available in the

literature, upon which this summary is based : Malka et al. (2008, 2013); Esarey et al. (2009); Tajima (2010) and Cakir et al. (2020).

The remainder of the section focuses on the mechanisms most relevant to laser-driven ion acceleration (LDIA), particularly TNSA, which remains the most widely studied and experimentally robust ion acceleration regime to date. Scaling laws and theoretical insights into Radiation Pressure Acceleration (RPA) are also discussed, along with complementary mechanisms such as Collisionless Shock Acceleration (CSA), Magnetic Vortex Acceleration (MVA) and Break-Out Afterburner (BOA), to provide a broader perspective on the diverse landscape of laser-based ion acceleration.

1.3.1 Electron acceleration

Among the various mechanisms for accelerating electrons using ultra-intense lasers, LWFA has emerged as the most prominent and widely studied. The following words introduce the key physical principles of LWFA and its operational conditions.

LWFA is a plasma-based acceleration mechanism that operates in an underdense plasma. An underdense plasma is a plasma in which the electron density n_e is lower than the critical density n_c corresponding to the frequency ω of an incident electromagnetic wave. This condition can be expressed as :

Equation 1.12 : Underdense plasma condition for LWFA

$$n_e < n_c = \frac{\varepsilon_0 m_e \omega^2}{e^2} \tag{1.12}$$

where n_e is the electron density of the plasma, n_c is the critical electron density, $\varepsilon_0 = 8.854 \times 10^{-12}$ F.m^{-1} is the vacuum permittivity, ω is the angular frequency of the laser ($\omega = 2\pi f$), m_e is the electron mass and $e = 1.6 \times 10^{-19}$ C. In this regime, the laser can propagate through the plasma, unlike in overdense media, where it is reflected. As an ultra-intense, ultra-short laser pulse travels through the underdense plasma, it exerts a ponderomotive force (Esarey et al., 2009) :

Equation 1.13 : Ponderomotive force exerted by an ultra-intense laser

$$\vec{F_p} = -m_e c^2 \vec{\nabla}\left(\frac{a_0^2}{2}\right) \tag{1.13}$$

where $a_0^2 = \frac{e^2}{m_e^2 c^2 \omega^2} |\vec{E}|^2$ is the square of the normalized vector potential of the laser and $c = 299792458$ m.s^{-1} is the speed of light. The ponderomotive force expels electrons from its path, creating a positively charged ion cavity. This displacement of charge launches a plasma wave (wakefield) trailing the pulse, oscillating at the plasma frequency ω_p. The plasma frequency determines the natural oscillation rate of electrons in the plasma, thereby setting the characteristic

wavelength and phase velocity of the excited wake. As detailed in the work of Esarey et al. (2009), the amplitude and phase velocity of the wakefield can be directly related to the plasma frequency :

Equation 1.14 : Plasma frequency

$$\omega_p = \sqrt{\frac{n_e e^2}{m_e \varepsilon_0}} \tag{1.14}$$

which determines the characteristic scale of the acceleration structure. In the linear regime, the corresponding electric field supported in the plasma can be expressed as :

Equation 1.15 : Maximum accelerating electric field in plasma (linear regime)

$$E_0 = \frac{c\, m_e\, \omega_p}{e} = \sqrt{\frac{n_e m_e c^2}{\varepsilon_0}} \tag{1.15}$$

Equation 1.15 shows that the maximal electric field sustainable in plasma depends on the electron density n_e since the remainder of the expression is constituted of fundamental constants. Conventional gas jets have densities on the order of $n_e \sim 10^{18}$ cm^{-3}, after insertion into equation 1.15 leads to $E_0 \sim 100$ GV/m, which is three orders of magnitude higher than the maximal electric field reachable in RF cavities, being limited by the cavity material's electrical breakdown. If electrons are injected into this wake either via self-injection (in the nonlinear bubble regime) or through controlled schemes. This acceleration occurs over the dephasing length of the plasma wake, which is usually on the order of several tens to hundreds of micrometers, determined by the Rayleigh length of the focal volume. Hence, this enables the acceleration of electron bunches to energies that are equivalently reached in conventional accelerators that are several tens of meters long, compared to the sub-millimetric scale with LWFA. Nowadays, gas nozzle technologies have undergone major developments, allowing them to reach electron densities of $n_e \sim 10^{20} - 10^{21}$ cm^{-3}, thus opening the path to generate accelerating gradients on the order of several TV/m.

Following the seminal paper of Toshiki Tajima and John Dawson in 1979 (Tajima et al., 1979), along with the advent of CPA-based lasers (Strickland et al., 1985) in 1985, several early experiments to perform all-optical electron acceleration occurred in the 1990s up to 2001. The resulting electron beams were characterized by ∼100 MeV maximal energy with Maxwellian energy spectrum (i.e. 100% energy spread exponentially decreasing with the vast majority in the lower energy section below 10 MeV) and large beam divergences (Kitagawa et al., 1992; Clayton et al., 1993; Everett et al., 1994; Modena et al., 1995; Nakajima et al., 1995; Umstadter et al., 1996; Ting et al., 1996; Moore et al., 1997; Amiranoff et al., 1998; Gahn et al., 1999; Malka et al., 2001). These characteristics were obtained by using Joule-level short-duration (tens of femtoseconds) laser pulses and focusing the laser energy in a sphere with a radius smaller than the plasma wavelength, which completely expels the electrons locally. This leads to the formation of a positively charged cavity

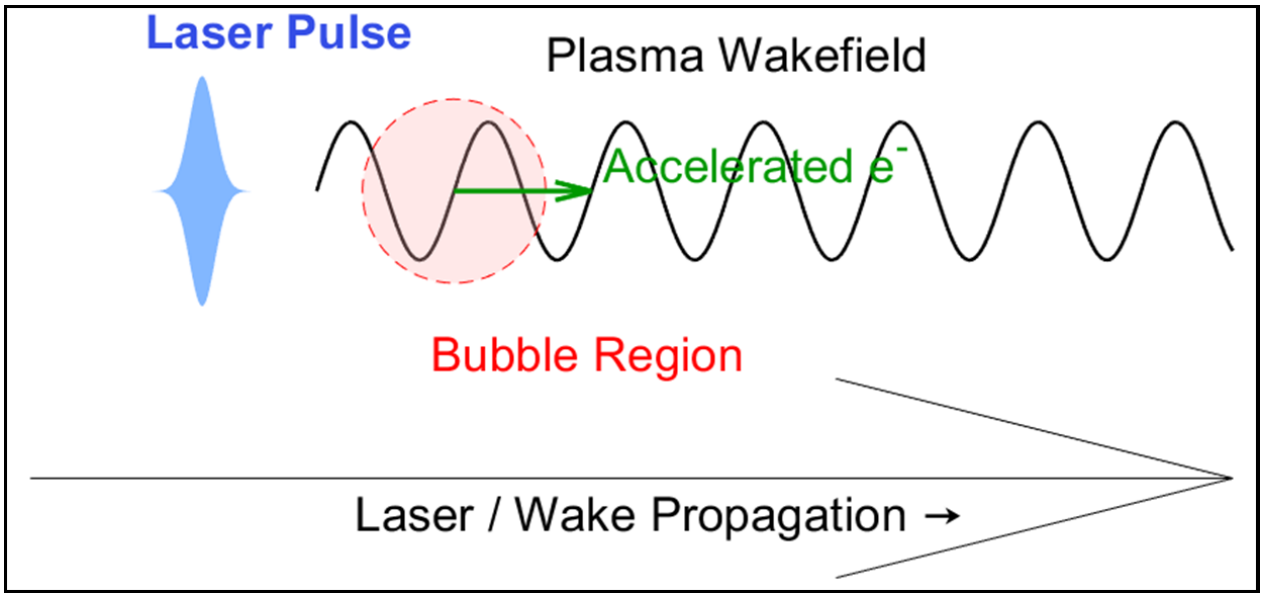

FIGURE 1.5 : Laser wakefield acceleration in the bubble regime.

The high-intensity laser pulses drive a plasma wave in its wake, which generates a longitudinal electric field that accelerates an electron bunch.

and thus to high electron trapping, a mechanism called the bubble or blow-out regime, depending on the laser intensity used (see Figure 1.5). The high-quality electron beams resulting from the bubble regime were rapidly perfected to reach 1 GeV electron energies in 2006 (Leemans et al., 2006)), and were further optimized over the recent years (Faure et al., 2006; Karsch et al., 2007; Hafz et al., 2008; Geddes et al., 2008; Froula et al., 2009; Clayton et al., 2010; Lu et al., 2011; Kim et al., 2013; Leemans et al., 2014; Kim et al., 2017). The record electron energy of nearly 8 GeV was reached in 2019 by the work of Gonsalves et al. (2019).

As described by Esarey et al. (2002) and Kostyukov et al. (2003), the strong ionization of the gaseous medium during LWFA creates an ion channel that exerts a restoring force on expelled electrons. While being accelerated longitudinally, the relativistic electrons undergo transverse betatron oscillations, leading to the emission of synchrotron-like X-ray radiation, known as betatron radiation. Thus, LWFA simultaneously acts as both an accelerator and a compact plasma wiggler, producing ultrashort (tens of femtoseconds), highly coherent X-rays with spectral properties similar to those from conventional undulators and wigglers. The first experimental observation of betatron radiation from LWFA was reported in 2004 by Rousse et al. (2004).

Insightful reviews on recent advances in betatron sources driven by LWFA can be found in the works of Albert et al. (2014) and Albert et al. (2016). One of the key advantages of laser-based betatron radiation (Figure 1.6.A) lies in its high spatial coherence, which enables the use of advanced imaging techniques such as X-ray Phase Contrast Imaging (XPCI) (Figures 1.6.B and 1.6.C). This technique, which leverages phase retrieval to significantly enhance image resolution, has become a routine diagnostic tool in LWFA experiments. Notably, the X-ray yield produced from a single

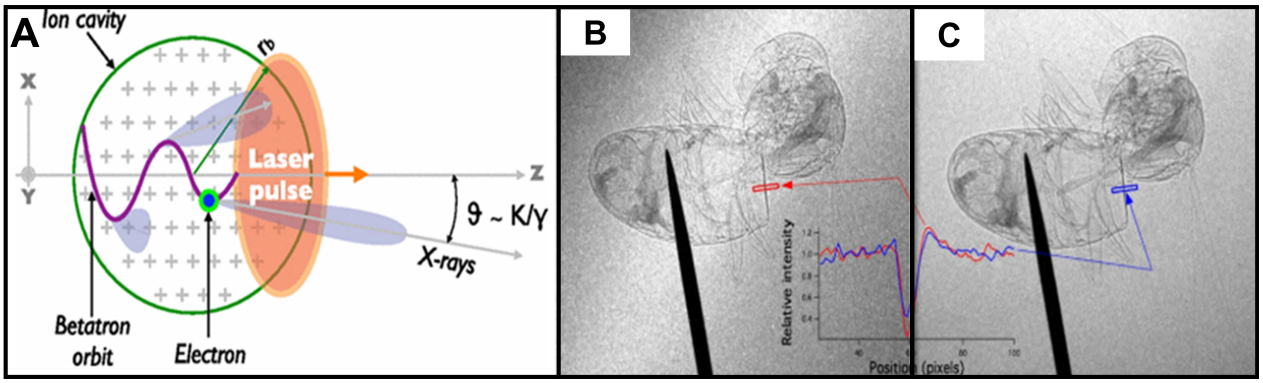

FIGURE 1.6 : Betatron x-ray mechanism and imaging.

(A) Schematic of the betatron mechanism. When an electron is injected into the ion cavity, it is submitted not only to the accelerating force but also to a restoring transverse force, resulting in its wiggling around the propagation axis. Because of this motion, the electron radiates x-rays whose typical divergence is $\theta = \kappa/\gamma$. (A) This figure has been adapted from Corde et al. (2013), with the publisher's permission. (B) and (C) are images of a bee with the x-ray Betatron beam with an edge line out indicated by the rectangular area : (B) one laser shot; (C) 13 laser shots. (B) and (C) have been adapted from Fourmaux et al. (2011a), with the publisher's permission. © Optica Publishing Group.

laser shot is not only competitive with synchrotron-based sources but also offers unique and complementary features. These capabilities have been demonstrated in several studies : Fourmaux et al. (2018); Vargas et al. (2019); Hussein et al. (2019); Fourmaux et al. (2020); Doherty et al. (2023) and Senthilkumaran et al. (2024).

1.3.2 Pioneering experiments and the establishment of Target Normal Sheath Acceleration

LDIA is a rapidly evolving field that leverages ultra-intense laser pulses ($> 10^{18}$ W.cm^{-2}) to accelerate ions to high energies via their interaction with matter. The field was launched by three landmark experiments in 2000, which reported the generation of multi-MeV proton beams from solid targets irradiated by high-intensity lasers (Snavely et al., 2000; Clark et al., 2000; Maksimchuk et al., 2000). These findings demonstrate the potential of compact, tabletop laser systems to rival and complement large-scale particle accelerators in terms of particle energies, opening the door to novel applications.

In one of the first studies, Clark et al. (2000) used a 1 ps laser pulse at intensities around 5×10^{19} W.cm^{-2} to irradiate aluminum foils and observed the emission of multi-MeV protons directed normal to the rear surface of the target. They attributed the acceleration to transient sheath electric fields generated by hot electrons transiting the target and escaping into vacuum, creating space-charge separation at the rear interface. This work provided early evidence of what would later be formalized as the TNSA mechanism. In parallel, Maksimchuk et al. (2000) irradiated thin aluminum foils with 400 fs laser pulses at similar intensities and observed forward-directed proton beams with

energies up to 1.5 MeV and narrow angular divergence. Like Clark et al. (2000), the protons were originating from the front target surface and crossed the full target thickness. Maksimchuk et al. (2000) also set up early experimental attention to the influence of ion species on acceleration. The most comprehensive and impactful demonstration came from Snavely et al. (2000), who employed the *Nova Petawatt* laser to deliver 500 fs pulses at a peak intensity of 3×10^{20} W.cm^{-2} onto gold foils. They reported the generation of proton beams reaching up to 58 MeV with high flux and narrow divergence, conclusively demonstrating rear-side acceleration consistent with a robust sheath field. With proper target preparation, heavier ions could also be accelerated to comparable energies per nucleon. The explanation came the following year from the work of Wilks et al. (2001), in a description of the mechanism called the TNSA, and received correct mathematical modeling of plasma expansion in vacuum in the work of Mora (2003). Crucially, Snavely et al. (2000) introduced and substantiated the TNSA mechanism, wherein relativistic laser-generated electrons form a quasistatic sheath field at the rear target surface, leading to efficient proton acceleration from surface contaminants. A schematic diagram of TNSA is shown in Figure 1.7. Supported by 2D particle-in-cell simulations, this framework rapidly became the foundation for understanding LDIA using solid targets. Collectively, these pioneering studies established the fundamental physics of laser-based ion acceleration, highlighted the importance of fast electron transport, target surface conditions, and sheath dynamics, and propelled the field toward the systematic investigation of laser-plasma coupling and optimization of beam parameters. The process also required sufficiently high laser-to-prepulse intensity contrast to avoid the pre-expansion of the target before the interaction with the main pulse, which would degrade the ion generation efficiency. Indeed, a nanosecond-long contrast of at least 10^{-6} was necessary when using intensities above 10^{18} W.cm^{-2}) since the ionization threshold for most atoms is around intensities of 10^{13} W.cm^{-2}. The ionization of atoms before the main pulse causes a charge imbalance followed by a plasma expansion, which disturbs the laser pulse propagation before reaching the interaction point with the target bulk. These prepulses originate primarily from the Amplified Spontaneous Emission (ASE) pedestal of the amplification chain and imperfections in pulse stretching or compression. Inconsistent explanations of the phenomenon driving the acceleration occurred among those publications.

1.3.3 Beam properties in Target Normal Sheath Acceleration

Building on these early demonstrations and armed with higher-contrast laser systems, better-diagnosed experiments, and refined theoretical models, the community soon shifted its focus from simply proving ion acceleration to mapping and engineering the full phase-space properties of the beams. Systematic studies of laser contrast, pulse duration, focal geometry, and target tailoring revealed clear pathways to higher energies, larger bunch charges, and tighter emittances, bringing TNSA sources into direct competition and complementarity with conventional accelerators. The next milestone, therefore, is not the mere observation of acceleration but the quantitative description of the beams that modern setups can now routinely produce.

These laser-driven ion beams are accelerated with kinetic energies up to 150 MeV for protons (Ziegler et al., 2024) and are surrounded by a co-moving electron cloud which forms a quasi-neutral plasma bunch that ensures beam stability and hence prevents its explosion due to electrostatic repulsion (i.e. the so-called Coulomb explosion). Ion beams also have short bunch duration (fs to ps at the source, three orders of magnitude shorter than what is achieved in conventional accelerators), can contain 10^{11} to 10^{13} protons per laser shot i.e. high bunch charges with a current in kA range if we remove the co-moving electrons, have a conical beam divergence of ($< 25^{\circ}$ half-angle), as well as an intrinsically low normalized transverse emittance as low as 4 μm.mrad (Cowan et al., 2004; Borghesi et al., 2004, 2008; Nürnberg et al., 2009) due to the small initial source size (a few microns) and short acceleration time. Reported beam divergences in TNSA experiments range from 5° to 25° full angle, with typical values around $\sim 10^{\circ} - 20^{\circ}$, depending on target thickness and laser focusing conditions.

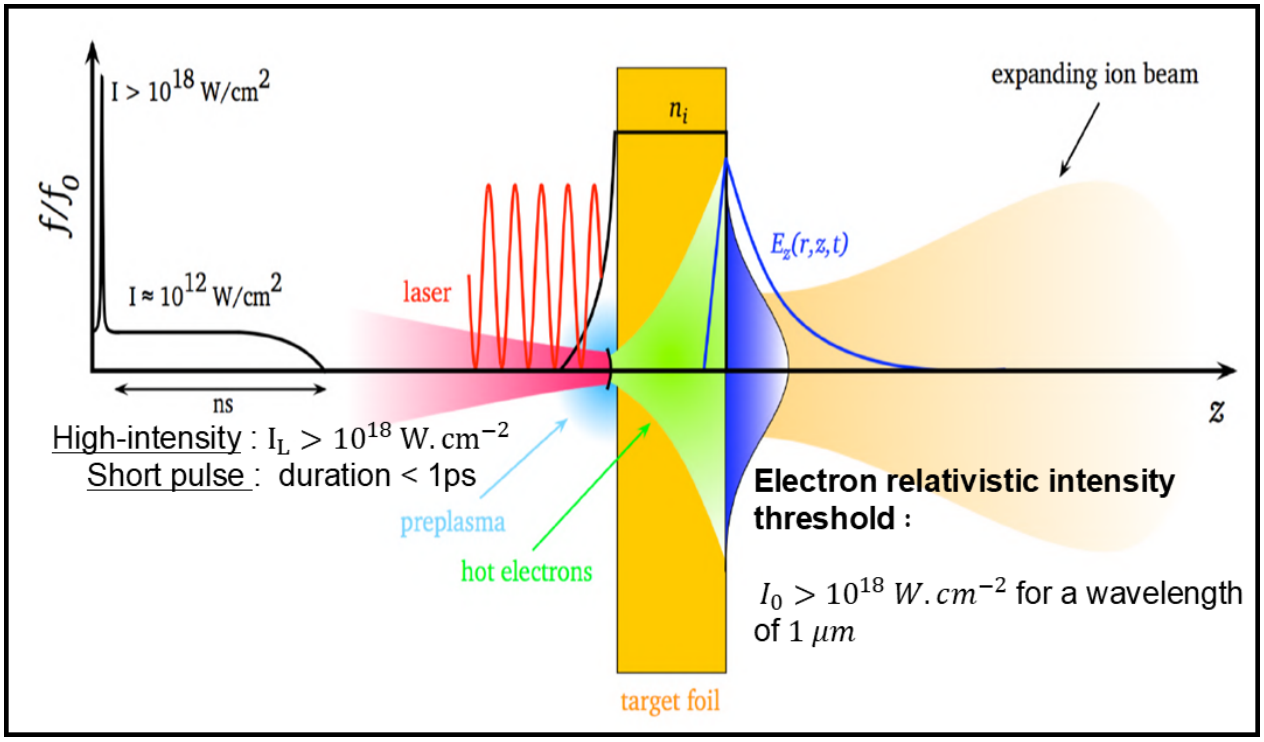

FIGURE 1.7 : Schematic of Target Normal Sheath Acceleration (TNSA).

Simplified representation of Target Normal Sheath Acceleration (TNSA) using solid targets. A high-intense short laser pulse interacts with the front surface of the target, generating hot electrons that propagate through the foil and form an intense sheath field at the rear surface, which accelerates ions to high energies. This figure has been adapted from Roth et al. (2016), with the publisher's permission.

1.3.4 Laser-plasma accelerators : A high-field alternative to RF-based systems

Although conventional accelerators are well-established and widely implemented, they are fundamentally constrained by the material breakdown threshold of their accelerating structures. This limitation caps the maximum achievable accelerating gradient to around 10-100 MV/m, depending on the cavity geometry and operating frequency. Exceeding these limits results in vacuum breakdown followed by an electrical discharge, which can damage components and degrade performance. Consequently, to reach particle energies in the hundreds of MeV or GeV range, conventional accelerators must span tens to thousands of meters, as seen in facilities like synchrotrons or linacs used in cancer therapy or high-energy physics. In addition, the equipment needed for beam transport and focusing - such as magnets and vacuum systems - makes conventional accelerators large, expensive, and complex. In contrast, laser-plasma accelerators use the collective behavior of plasma to generate electric fields above 1 TV/m, over 1,000 times stronger than those in traditional RF accelerators. These extreme fields arise from ultrashort (<1 ps), high-intensity laser pulses ($\geq 10^{18}$ W.cm^{-2}) that generate highly relativistic electrons and drive strong space-charge separation. The resulting electrostatic or radiation pressure fields can efficiently accelerate electrons over millimeter to centimeter distances and ions over sub-micrometer to millimeter scales, depending on the regime of interaction. Among the various ion species, protons are preferentially accelerated due to their high charge-to-mass ratio ($Z/A = 1$), which makes them more responsive to the rapidly evolving sheath fields generated in TNSA (Macchi et al., 2013). Their low mass allows them to reach high velocities within the short acceleration timescale, and they can follow temporal variations of the accelerating field more closely than heavier ions (Daido et al., 2012). Additionally, protons often originate from contaminant layers (e.g., hydrocarbons or water) on the target rear surface, placing them directly at the peak of the sheath field and enhancing their acceleration efficiency (Toncian et al., 2006). Unless these layers are removed, heavier ions typically exhibit lower energies due to their smaller charge-to-mass ratio and delayed response to the sheath field (Hegelich et al., 2006).

1.3.5 Foundational reviews in LDIA

These previous observations are further supported by three influential reviews that have shaped the field. In particular, the comprehensive work by Macchi et al. (2013) classifies LDIA mechanisms based on the laser intensity, target thickness, and ion species involved, with particular emphasis on TNSA, the most widely studied regime to date. This study also examines alternative mechanisms such as RPA and BOA, which become dominant at higher laser intensities and with ultrathin or transparent targets. The review discusses :

— Scaling laws for maximum ion energy and ion number,

— Effects of laser pulse contrast, focal geometry, and target structuring,

— Instabilities that can degrade beam quality,

— Developments in diagnostics and beam characterization,

— Applications in proton therapy, nuclear physics, and ultrafast probing.

Finally, the authors stress the importance of numerical simulations, particularly Particle-In-Cell methods, to predict experimental outcomes and explore unexplored regimes. The paper concludes with an outlook on challenges and prospects, including high-repetition-rate operation, beam transport, and laser technology improvements required for real-world applications. Daido et al. (2012) provide a comprehensive review outlining analytical models based on plasma expansion into vacuum, especially those developed by Mora (2003), which capture the scaling of maximum ion energy with laser and target parameters. These models relate the ion energy to the hot electron temperature, itself dependent on the laser intensity through the ponderomotive scaling. The paper also discusses refinements accounting for finite sheath sizes and electron divergence, which are critical for matching experimental data. TNSA is also contrasted with other mechanisms such as Coulomb explosion, RPA, and CSA, but it stands out for its relative simplicity, reproducibility, and compatibility with currently available TW- and PW-class laser systems. Experimental validation of the TNSA scaling laws across multiple facilities, as well as its applications in proton radiography, isochoric heating, and medical therapy, underscores its versatility and practical relevance. More recently, the book *Applications of Laser-Driven Particle Acceleration*, edited by Bolton et al. (2018), presents an overview of both the fundamental principles and the emerging applications of laser-driven particle sources. Divided into two parts, the first section addresses the underlying physics and technical capabilities of electron and ion acceleration using high-intensity lasers, including associated X-ray, γ-ray, and neutron generation. The second part explores a broad spectrum of envisioned applications, such as ion beam radiotherapy, space radiation studies, radiobiology, nuclear reactions in plasmas, radioisotope production, materials analysis, and advanced imaging techniques. A central concept emphasized throughout the book is the "uniqueness strategy" which advocates for leveraging the distinctive features of laser-driven particle beams, such as ultrashort pulse duration, compactness, and multi-species synchronization, to complement and extend the capabilities of conventional accelerator systems. Indeed, authors propose that laser-based accelerators should not try to replace conventional particle accelerators but instead offer new capabilities like very short bursts of particles or multiple types of particles at once, which conventional accelerators can't provide. This perspective reinforces the potential of laser-based sources to contribute meaningfully to the advancement of accelerator-driven applications.

1.3.6 Physics of Target-Normal Sheath Acceleration

The Target-Normal Sheath Acceleration mechanism occurs when a high-intensity ($I_L \geq 10^{18}$ W.cm^{-2}), short-duration (≤ 1 ps) laser pulse interacts with a thin solid target, typically of micrometric thickness. For efficient acceleration, it is essential to maintain an unperturbed target surface up to the arrival of the main pulse. This requires a high temporal contrast ratio ($\geq 10^6$) to suppress prepulses that might induce premature plasma formation. Above the electron's relativistic intensity

threshold ($I_L \geq 10^{18}$ W.cm^{-2} for a laser wavelength of 1 μm), the intense EM wave pushes the electrons from the target, which are easily ionized mainly through high-field tunnel ionization, in the laser's forward direction through the relativistic ponderomotive force. The accelerated electrons undergo strong heating mostly through $\vec{J} \times \vec{B}$ heating and vacuum heating absorption mechanisms in the intensity range of [10^{18};10^{20}] W.cm^{-2}. Hot electrons then propagate through the target at relativistic velocities and exit the target's rear side, where some of them are retained due to strong charge separation, based on the target's capacitance, and form an electron cloud (the "sheath") that establishes an intense quasi-electrostatic electric field over a distance on the order of the Debye length λ_D. The longitudinal electric field completely ionizes water molecules and hydrocarbon impurities instantaneously, which are naturally present on the surface, and ions get accelerated along the target-normal direction. The acceleration terminates when both the ejected ion and electron populations have reached the same thermal velocities to form a quasi-neutral plasma bunch. Different ion species can be accelerated by this mechanism; however, protons are of particular interest since they respond more quickly to the accelerating electric field, which results in stronger energy absorption from the hot electron cloud, owing to their higher charge-to-mass ratio compared to carbon ion species, for instance. The same plasma expansion also occurs on the target's front surface; however, the acceleration is more efficient from the rear surface in the forward direction as a result of sharper density gradients, leading to stronger electric fields. More precisely, the front surface's plasma gradient is more affected by the EM wave's perturbation that produces a blow-off plasma, which hinders the conservation of a steep plasma density gradient. The high intensity of the EM wave, combined with the fact that ions start at rest on a cold surface, leads to strongly forward-oriented (i.e., non-isotropic) and relatively low divergence and laminar energetic ion beams, hence providing high beam quality. The short laser pulse duration ensures the generation of a steep hot electron density gradient and thus of a strong accelerating electric field. Another crucial factor is the laser contrast, which must be high enough to prevent a too-strong target decompression before the interaction with the main pulse. The TNSA mechanism, therefore, relies on the generation of a dense hot electron cloud as an intermediate step in the laser-to-proton energy transfer. The acceleration process (Wilks et al., 2001) can be modeled by a plasma expansion model (Mora, 2003).

1.3.7 Electron heating and sheath formation : Electromagnetic framework of Target-Normal Sheath Acceleration

Maxwell's equations in differential form provide a fundamental basis for describing the field dynamics governing TNSA. In this study, we explore how each equation contributes to the different stages of electron heating, sheath formation, and ion acceleration. Those different stages are shown in Figure 1.8.

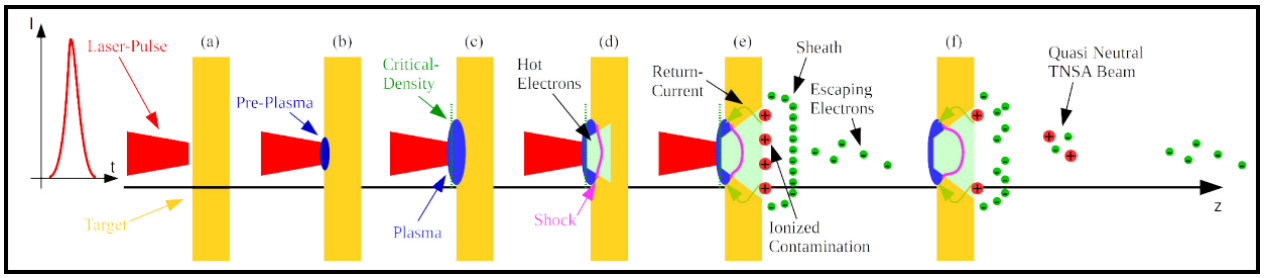

FIGURE 1.8 : Stages of TNSA-driven proton acceleration.

Illustration of the key stages leading to Target Normal Sheath Acceleration (TNSA) of protons. (A) a laser with finite contrast interacting with an initially solid target. The rising edge of laser intensity creates a pre-plasma (B) and heats it. Plasma density rises up to a critical density (C) where the intensity peak becomes reflected. A shock wave is created (D), and meanwhile electrons become forward-accelerated. Escaping fast electrons build up a potential that stops slower electrons, which form a sheath (E). Contamination on the target rear side is ionized and subsequently accelerated (F). Accelerated ions take electrons with them, forming a quasi-neutral TNSA beam. The process ends due to a disturbance by the shock or by the end of laser-electron acceleration. With time, fewer electrons are within the sheath, and the acceleration efficiency drops. This figure has been adapted from Ehret (2016), with the author's permission.

Electron heating and laser-plasma interaction in Target Normal Sheath Acceleration

The initial stage of TNSA is dominated by the interaction of an ultra-intense laser pulse with a solid target (Figure 1.8.A). The rising edge of the laser intensity generates and heats a pre-plasma at the target front (Figure 1.8.B). The plasma density grows until it reaches the critical density, where the main intensity peak is reflected (Figure 1.8.C). This interaction launches a shock wave, while electrons are driven forward into the target (Figure 1.8.D). The fastest electrons escape, creating an electrostatic potential that traps slower electrons and forms a sheath at the target rear (Figure 1.8.E). Understanding the mechanisms governing electron heating and transport is therefore essential for optimizing ion acceleration. At laser intensities exceeding 10^{18} W.cm^{-2}, relativistic effects dominate the electron dynamics, leading to various absorption mechanisms that transfer laser energy to the electrons. With increasing laser intensity, the prevailing electron-heating mechanisms transition through several regimes, including :

1. **Resonant absorption** (Wilks et al., 1992) : Occurs under p-polarization and requires a longer density gradient at the critical surface to match the laser frequency to the local plasma oscillation.

2. **Vacuum heating (Brunel absorption)** (Brunel, 1987) : Efficient for p-polarized light at oblique incidence, particularly in steep density gradient plasmas.

3. $\vec{J} \times \vec{B}$ **heating** (Kruer et al., 1985) : A relativistic mechanism dominant at normal incidence and high intensities, particularly for linearly polarized pulses (effective for both s- and p-polarization, and even at normal incidence).

4. **Ponderomotive acceleration** (Wilks et al., 1992) : Arises from the inhomogeneous laser field pushing electrons away from high-intensity regions; relevant under most conditions but

especially in the relativistic regime. Works with both p- and s-polarization. No requirement for incidence angle : can be normal or oblique.

These mechanisms govern the coupling of laser energy into relativistic electrons, which in turn drives plasma expansion and ion acceleration. Each of these mechanisms can be analyzed through Maxwell's equations in differential form.

Ampère's Law (equation 1.8) describes the generation of magnetic fields by the hot electron current. As a high-intensity laser pulse strikes a solid target, it drives forward-propagating hot electron currents through the plasma. To maintain charge neutrality, return currents flow in the opposite direction within the colder background plasma. These opposing current systems create strong azimuthal and time-variant magnetic fields via Ampère's Law. The resulting self-consistent magnetic fields significantly influence the transport and confinement of hot electrons, thereby shaping the spatial and temporal characteristics of energy deposition at the target rear surface.

Faraday's Law (equation 1.9) describes how time-varying magnetic fields induce electric fields. The temporal evolution of those magnetic fields gives rise to azimuthal electric fields, which in turn influence the confinement and transport of hot electrons. This feedback mechanism affects both the spatial localization of electron energy deposition and the efficiency of sheath field formation.

Gauss's Law (equation 1.10) relates the electric field to the local charge density, and thus governs charge separation. As the laser drives hot electrons forward, many escape the target surface, leaving behind a region of net positive charge. This localized charge imbalance generates strong electrostatic fields, which in turn influence the dynamics of electron escape and sheath formation. These quasi-static sheath fields directly accelerate surface ions. Consequently, Gauss's Law underpins the spatial structure and strength of the accelerating field, and thereby strongly impacts ion acceleration efficiency.

Mechanisms of laser energy absorption by electrons

Several mechanisms contribute to electron heating in TNSA, each dependent on laser intensity, polarization, and angle of incidence.

Resonant absorption is a collisionless energy transfer mechanism that occurs when a p-polarized electromagnetic wave impinges obliquely on a plasma with a steep density gradient. Under these conditions, the laser frequency matches the local plasma frequency near the critical surface, enabling efficient coupling of laser energy into plasma oscillations. As shown schematically in Figure 1.9, this resonance condition leads to the excitation of electron plasma waves, which grow rapidly during the initial optical cycles of the laser pulse. The absorbed energy is subsequently transferred to plasma electrons through mechanisms such as wave damping from charged particle collisions, Landau damping, and wave breaking (Forslund et al., 1977; Estabrook et al., 1978). The plasma

frequency is given by Equation 1.14. The degree of absorption is dependent on the plasma density and the laser's wavelength. It represents the fraction of laser energy absorbed by the plasma following the relationship :

Equation 1.16 : Fraction of laser energy absorbed by the plasma (resonant absorption)

$$A \approx \frac{\pi}{2}\left(\frac{L_s}{\lambda_L}\right)^{1/3}\sin^2\theta \tag{1.16}$$

where λ_L is the laser wavelength, L_s is the plasma gradient scale length and θ is the angle of incidence of the laser.

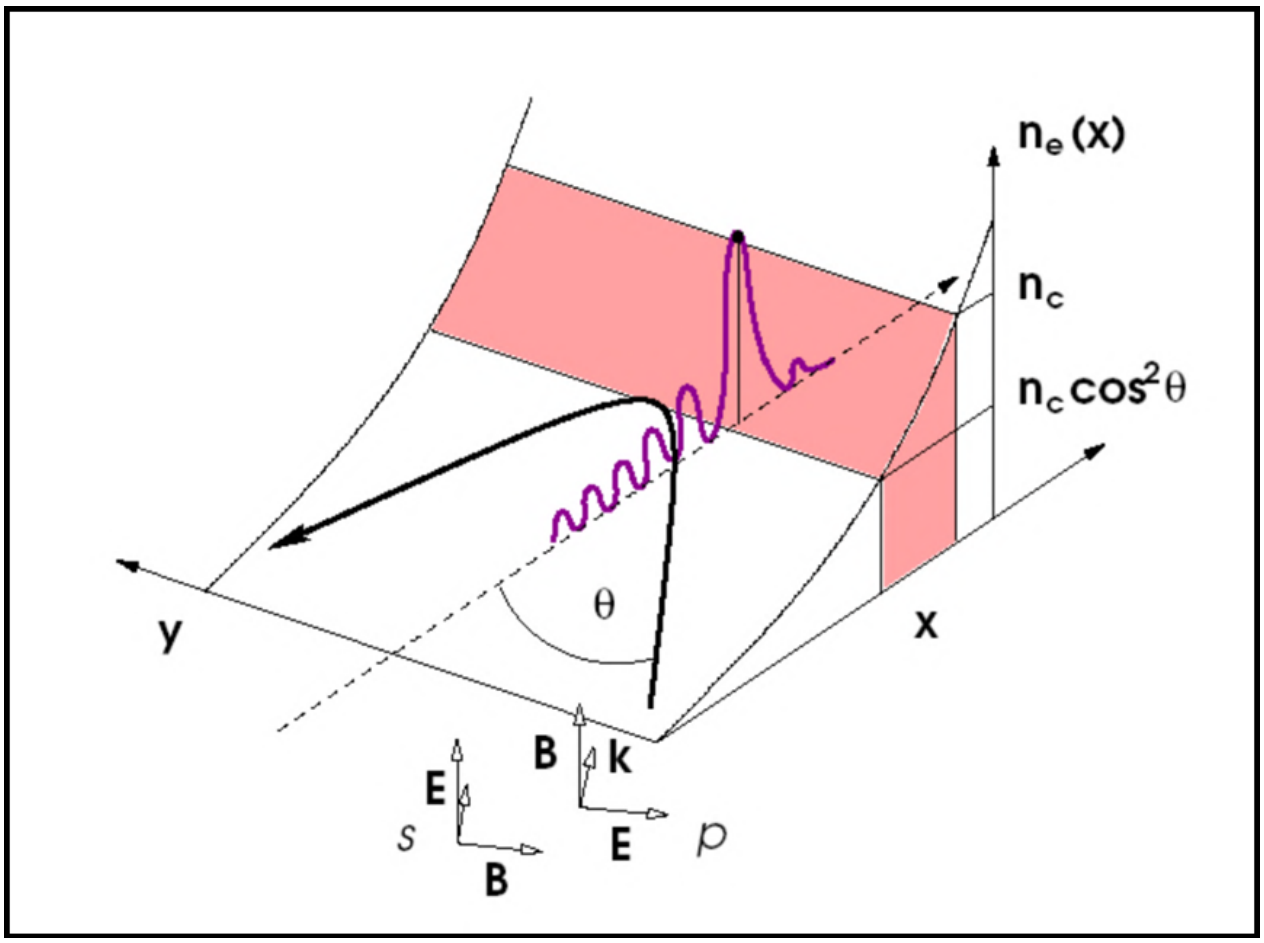

FIGURE 1.9 : Schematic of resonance absorption.

Schematic diagram of resonance absorption. Both linear polarizations are defined, the p-polarization being when the electric field is oriented in the EM wave's plane of incidence (yz-plane here), whereas the s-polarization is for an electric field oriented perpendicularly to the plane of incidence. This figure has been adapted from Gibbon (2005), with the publisher's permission.

It is important to note that resonant absorption contributes only marginally to electron heating as the laser intensity approaches the relativistic threshold, typically around 10^{18} W.cm^{-2}, beyond which its efficiency rapidly diminishes. Consequently, it does not constitute the dominant absorption mechanism in the relativistic regime relevant to TNSA. Resonant absorption is therefore not the dominant mechanism driving hot electron generation in TNSA. Nonetheless, resonant absorption can still contribute to initial electron heating during the rising edge of the laser pulse, especially when the intensity is still below the relativistic threshold. Its role is thus relevant in shaping the early-time hot electron population and influencing pre-plasma dynamics.

When a p-polarized EM wave with oblique incidence encounters a very steep plasma density gradient (i.e. $L_s/\lambda_L \leq 0.08$), resonant absorption becomes "not-so-resonant" since the sharp gradient hinders the growth of plasma waves. In this context, Brunel (1987) proposed an alternative and highly efficient absorption mechanism, now widely known as vacuum heating or Brunel absorption. Vacuum heating occurs when electrons are pulled from the plasma by the laser into the vacuum by the oscillating laser electric field before being reinjected with increased energy. This mechanism is dominant when the plasma density is well above the critical density and when the density gradient scale length is much shorter than the electron quiver length ($L_s \ll v_{osc}/\omega$), making it a key process in high-power laser-solid target interactions. Brunel absorption is one of the major absorption mechanisms for intensities in the range 10^{17} to 10^{19} W.cm^{-2}, above which the $\vec{J}\times\vec{B}$ heating absorption mechanism dominates. The Brunel absorption coefficient can reach ∼10% around the electron relativistic regime intensity threshold (laser intensity of $I_L = 1.37 \times 10^{18}$ W.cm^{-2} for $\lambda_L = 1$ μm). A particularly noteworthy feature of this mechanism is the generation of ultrashort electron bunches, emitted at the laser frequency $f_L = \omega/2\pi$, resulting in a periodic injection of high-energy electrons at every optical cycle. The energy absorbed per cycle can be estimated using the relation :

Equation 1.17 : Energy absorbed per optical cycle in Brunel absorption

$$K = eE_0L_s \tag{1.17}$$

where E_0 is the peak laser field strength.

Above the electron relativistic regime threshold, the oscillation velocity of electrons in the EM field approaches the speed of light, and the $\vec{v} \times \vec{B}$ term becomes of preponderant importance in the Lorentz force. $\vec{J} \times \vec{B}$ heating is effective for both linear polarizations s and p, however, completely suppressed for circular polarization. The intense electric field first makes the electrons oscillate from the ponderomotive force in the transverse plane, providing them with a relativistic velocity and thus a strong oscillatory electron current $\vec{J_e}$. The steadily growing current density $\vec{J_e}$ then concomitantly interacts with the B-field of the EM wave and generates a Lorentz force that pulls the electrons out of the target to finally re-inject them back in. Since $\vec{v}$ and $\vec{B}$ both oscillate at the same

frequency, the resulting term oscillates at the doubled frequency and therefore re-injects electron bunches in the target twice per laser period. The relativistic electrons that acquire kinetic energy on the order of the ponderomotive potential go far beyond the skin depth and Debye length, away from plasma influence, which makes them shielded from the laser field thereafter. This absorption mechanism occurs similarly to vacuum heating with the presence of steep-edged density gradients and high intensities; however, in the $\vec{J} \times \vec{B}$ mechanism, the B-field is non-negligible and pushes the relativistic electrons along the longitudinal axis even for s-polarization, being most effective for normal incidence. It is the combination of both the effects of electrons oscillating at relativistic velocities and the non-negligible B-field in the Lorentz force that gives rise to the $\vec{J} \times \vec{B}$ heating as a non-linear ponderomotive force. This absorption process dominates the TNSA mechanism for intensities in the range 10^{18} to 10^{20} W.cm^{-2}, and is explained in Kruer et al. (1985). A few years later, Wilks et al. (1992) has shown through simulations that the average kinetic energy of the relativistic electrons re-entering the target, also called hot electron temperature $k_B T_{hot}$ can be properly estimated by the relativistic ponderomotive energy :

Equation 1.18 : Relativistic hot-electron temperature from ponderomotive scaling

$$k_B T_{hot} \approx m_e c^2 \left(\sqrt{1 + a_0^2} - 1 \right) \tag{1.18}$$

where $a_0 = \frac{eE_0}{m_e \omega c}$ is the normalized laser vector potential, directly influencing the energy of the laser-driven ions. The hot electrons escaping into the vacuum establish a charge separation, leading to the formation of a strong sheath field (∼TV/m), which accelerates ions normal to the target surface. Efficient $\vec{J} \times \vec{B}$ heating enhances the sheath field in TNSA experiments. It is important to note here that the ponderomotive energy is a reasonable estimate for the electron temperature and provides the proper scaling, but does not provide a rigorous description of electron heating. The absorption coefficient $A_{\vec{J} \times \vec{B}}$ can be written as follows :

Equation 1.19 : Absorption coefficient for $\vec{J} \times \vec{B}$ heating

$$A_{\vec{J} \times \vec{B}} \sim \frac{1}{2\pi} \left[1 - \left(1 + a_0^2 \right)^{-1/2} \right] \tag{1.19}$$

Similarly to the Vacuum heating, the $\vec{J} \times \vec{B}$ heating absorption coefficient will reach an asymptotic value with increasing laser intensity, going up to 16% for high values of a_0 as shown in Figure 1.10. The Brunel absorption coefficient describes the direct vacuum heating of electrons at steep plasma–vacuum interfaces, where oscillating electrons are periodically pulled into and out of the plasma by the laser electric field. For oblique incidence, it can be approximated by :

Equation 1.20 : Brunel absorption coefficient for oblique incidence

$$A_B \approx 0.05\, a_0^{-1} \tan\theta \left[\sqrt{1 + a_0^2 \sin^2\theta} - 1\right] \tag{1.20}$$

where a_0 is the normalized vector potential and θ the incidence angle. The saturation observed in Figure 1.10 originates from the transition between the sub-relativistic and relativistic regimes of the $\mathbf{J} \times \mathbf{B}$ heating process. At moderate intensities ($a_0 \sim 1$), the electron's transverse velocity oscillates sinusoidally, and the longitudinal Lorentz force $F_z \sim v_\perp B_\perp$ produces asymmetric kicks that lead to a net forward drift and efficient absorption. However, as $a_0 \gg 1$, the electron motion becomes ultra-relativistic : its transverse oscillation flattens near $\pm c$, the Lorentz factor γ remains nearly constant, and forward and backward kicks become symmetric, causing the cycle-averaged current $\langle J \rangle \to 0$. Consequently, the absorption coefficient saturates at its theoretical upper limit of approximately 16% ($\simeq 1/2\pi$). In practice, this limit is rarely achieved because real laser–solid interactions deviate from the idealized model : finite contrast pulses generate preplasmas; multidimensional effects, surface roughness, and self-generated fields break the symmetry; competing mechanisms such as parametric instabilities and sheath formation divert part of the absorbed energy; and finally, collisions and return currents redistribute the absorbed energy within the target, preventing the clean saturation predicted by theory.

The transport properties of hot electrons play a crucial role in defining the efficiency of sheath formation. Studies using high-intensity laser pulses confirm that $\vec{J} \times \vec{B}$ heating significantly contributes to hot electron production, leading to proton energies exceeding tens of MeV, making it a critical process for optimizing LDIA. Hot electron temperatures of 1-5 MeV have been observed in high-intensity laser experiments (Fuchs et al., 2006), with fast electron transport distances reaching 100s of microns in experimental setups at *LULI*, *Vulcan*, and *Titan* laser facilities. PIC simulations confirm the dominance of $\vec{J} \times \vec{B}$ heating in the relativistic regime (Sentoku et al., 2008). These simulations further reveal the formation of return current-driven magnetic fields, which can reach strengths of 100 T (Gremillet et al., 2002), affecting electron transport and sheath formation.

Sheath Formation and Charge Separation

As foreseen in the previous section, when an ultra-intense laser pulse irradiates a solid target, energetic electrons are generated at the front surface and rapidly transported to the rear side of the target. A fraction of these hot electrons escapes into vacuum, leaving behind a positively charged ion layer. This results in a charge separation that gives rise to a quasi-static electrostatic sheath field normal to the target surface. The strength of this field can reach several TV/m and is among the most intense electric fields achievable in laboratory conditions (Wilks et al., 2001; Mora, 2003; Fuchs et al., 2006).

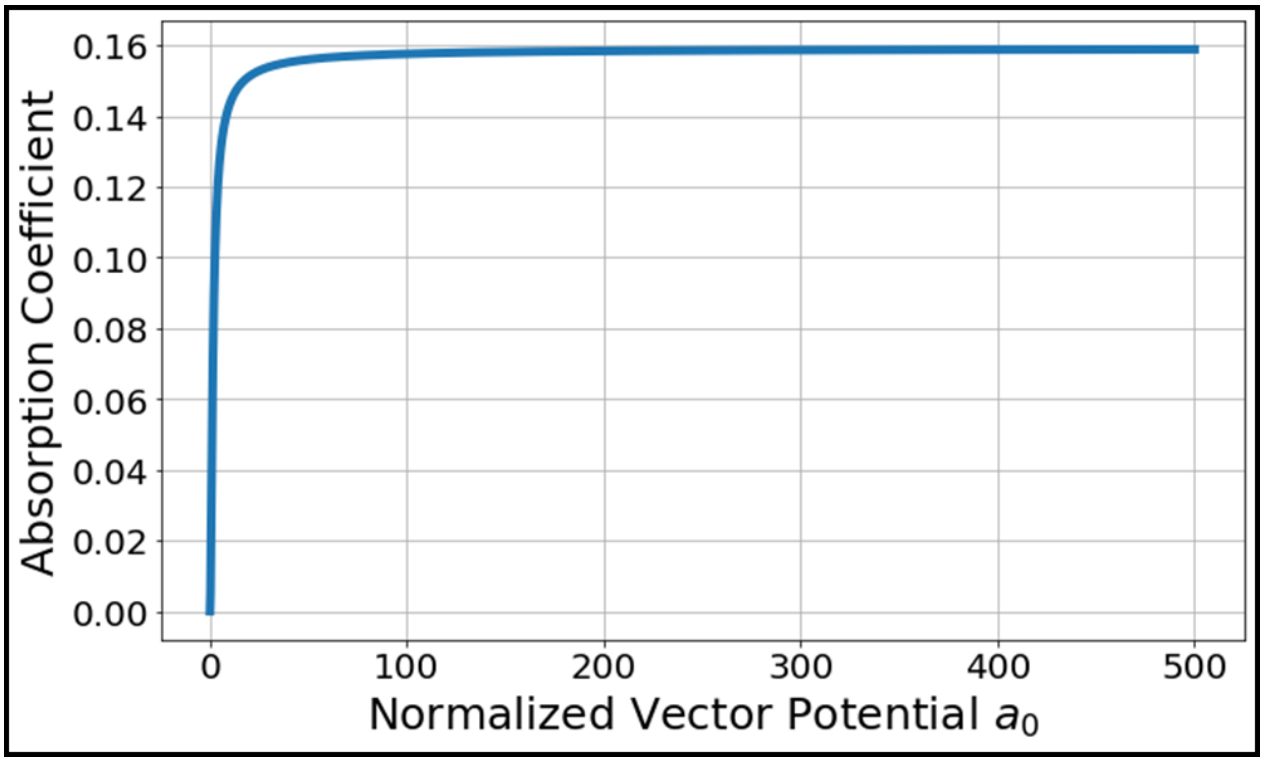

FIGURE 1.10 : Absorption saturation in the $\vec{J} \times \vec{B}$ regime.

$\vec{J} \times \vec{B}$ absorption coefficient as a function of the normalized laser vector potential a_0. The absorption increases rapidly in the sub-relativistic regime and asymptotically saturates near $1/(2\pi) \approx 0.159$ for $a_0 \gg 1$. This reflects the nonlinear saturation of energy transfer from the laser field to electrons in steep density gradient plasmas.

The spatial extent of the sheath is on the order of the Debye length :

Equation 1.21 : Debye length at the target front surface

$$\lambda_D(t) = \left(\frac{\varepsilon_0 k_B T_{hot}}{n_{e,\text{front}}(t)\, e^2} \right)^{1/2} \tag{1.21}$$

where T_{hot} is the hot electron temperature and $n_{e,\text{front}}(t)$ the time-dependent electron density at the sheath front, which evolves as :

Equation 1.22 : Time-dependent electron density at the sheath front

$$n_{e,\text{front}}(t) \simeq \frac{2\, n_{e0}}{\omega_{pi}^2\, t^2} \tag{1.22}$$

with n_{e0} the hot electron density and $\omega_{pi} = \sqrt{Z^2 e^2 n_i / m_i \varepsilon_0}$ the ion plasma frequency (Mora, 2003). It can be estimated from energy conservation and geometric considerations. The total number of hot electrons generated by the laser and transported to the rear surface is given by :

Equation 1.23 : Total number of hot electrons transported to the rear surface

$$N_e = \frac{f\,\mathcal{E}_L}{T_{\rm hot}} \tag{1.23}$$

where f is the fraction of laser energy converted into hot electrons, $\mathcal{E}_L$ is the total laser energy deposited on the target, and T_{hot} is the mean hot electron temperature. These electrons are injected over the laser pulse duration τ_{laser} and are assumed to populate a sheath region of area S_{sheath}. The corresponding hot electron density is then :

Equation 1.24 : Hot electron density at the sheath front

$$n_{e0} = \frac{N_e}{c\,\tau_{\rm laser}\,S_{\rm sheath}} = \frac{f\,\mathcal{E}_L}{T_{\rm hot}\,c\,\tau_{\rm laser}\,S_{\rm sheath}} \tag{1.24}$$

The sheath area accounts for the divergence of electrons within the target and is given by :

Equation 1.25 : Electron sheath area

$$S_{\rm sheath} = \pi\,(r_0 + d\,\tan\theta)^2 \tag{1.25}$$

where r_0 is the laser spot radius on the front surface, d is the target thickness, and θ is the typical divergence angle of the electrons inside the target ($\theta \sim 25°$, Coury et al. (2013)). In TNSA experiments, this yields electron densities in the range $n_{e0} \sim 10^{19}$–10^{21} cm^{-3}, depending on the laser intensity, conversion efficiency, and target geometry. Assuming typical TNSA conditions using *ALLS* 150 TW laser characteristics :

— Laser-to-hot-electron conversion efficiency : $f = 0.005$,

— Laser pulse energy : $\mathcal{E}_L = 3.2$ J,

— Laser pulse duration (FWHM) : $\tau_{laser} = 22$ fs $= 22 \times 10^{-15}$ s,

— Speed of light : $c = 3.00 \times 10^8$ m $\cdot$ s^{-1},

— Hot electron temperature : $T_{hot} = 1$ MeV $= 1.602 \times 10^{-13}$ J,

— Initial hot electron density : $n_{e0} = 1.0 \times 10^{27}$ m^{-3},

— Ion charge : $Z = 1$,

— Ion mass (protons) : $m_i = 1.67 \times 10^{-27}$ kg,

— Elementary charge : $e = 1.602 \times 10^{-19}$ C,

— Vacuum permittivity : $\varepsilon_0 = 8.854 \times 10^{-12}$ F $\cdot$ m^{-1},

Equation 1.24 reduces to $S_{\rm sheath} \simeq 1.51 \times 10^{-11}$ m^2. We then compute the ion plasma frequency :

$$\omega_{pi} = \sqrt{\frac{(1.602 \times 10^{-19})^2 \cdot 10^{27}}{1.67 \times 10^{-27} \cdot 8.854 \times 10^{-12}}} \approx 4.17 \times 10^{13}\ \mathrm{rad/s},$$

allowing to estimate the sheath-front density at $t = 1$ ps :

$$n_{e,\text{front}}(t) \simeq \frac{2 \cdot 10^{26}}{(4.17 \times 10^{13})^2 \cdot (1 \times 10^{-12})^2} \approx 9.2 \times 10^{26}\ \text{m}^{-3}.$$

Finally, since $k_B T_{hot} = 1.602 \times 10^{-13}$ J, we compute :

$$\lambda_D \approx \sqrt{\frac{8.854 \times 10^{-12} \cdot 1.602 \times 10^{-13}}{9.2 \times 10^{26} \cdot (1.602 \times 10^{-19})^2}} \approx 0.23\ \mu\text{m}.$$

Having estimated the key plasma parameters governing the sheath dynamics, we now introduce the fundamental equation that describes the electrostatic potential structure within the sheath region. This is given by Poisson's equation, which relates the spatial distribution of the potential to the local charge density :

Equation 1.26 : Poisson's equation

$$\nabla^2 \Phi = -\frac{\rho}{\varepsilon_0} \tag{1.26}$$

This equation forms the foundation for modeling the electrostatic potential Φ that arises from the local charge imbalance in the sheath formed at the rear side of a laser-irradiated target. The solution of this equation provides insights into the sheath expansion dynamics and the potential available for ion acceleration (Macchi et al., 2013). $\rho = e(Zn_i - n_e)$, where n_e is the hot electron density and n_i is the ion density. Assuming a Boltzmann distribution for the hot electrons (Mora, 2003), the density profile takes the form :

Equation 1.27 : Boltzmann density profile for hot electrons

$$n_e(x,t) = n_{e0} \exp\left(\frac{e\Phi(x,t)}{k_B T_{hot}}\right) \tag{1.27}$$

leading to a nonlinear form of Poisson's equation :

Equation 1.28 : Nonlinear form of Poisson's equation in the sheath region

$$\frac{d^2\Phi}{dx^2} = -\frac{e}{\varepsilon_0}\left[Z\, n_i(x,t) - n_{e0} \exp\left(\frac{e\,\Phi(x,t)}{k_B T_{\text{hot}}}\right)\right] \tag{1.28}$$

Solving this equation in combination with ion motion equations allows one to derive the sheath field and its temporal evolution. The electric field at the ion front, as derived by Mora (2003), evolves as :

Equation 1.29 : Temporal evolution of the electric field at the ion front (Mora model)

$$E_{\text{front}}(t) \simeq \frac{2\,E_0}{\sqrt{2e}\,\omega_{pi}\,t} \tag{1.29}$$

with $E_0 = \sqrt{\frac{n_{e0}k_BT_{hot}}{\varepsilon_0}}$ and $\omega_{pi} = \sqrt{\frac{Z^2e^2n_{i0}}{m_i\varepsilon_0}}$.

In TNSA, the maximum (cutoff) energy acquired by ions during rear-surface sheath expansion can be modeled using a time-limited, isothermal, self-similar fluid approach (Fuchs et al., 2006). The model assumes that a burst of hot electrons rapidly sets up a sheath field, and ion acceleration occurs over a limited time t_{acc}, approximately proportional to the laser pulse duration. The expression for the maximum ion energy is given by :

Equation 1.30 : Maximum ion energy from isothermal sheath expansion (Mora/Fuchs model)

$$\mathcal{E}_{max} = 2Z\,k_BT_{hot}\left[\ln\left(t_p + \sqrt{t_p^2+1}\right)\right]^2 \tag{1.30}$$

where T_{hot} is the hot electron temperature and $t_p = \omega_{pi}t_{\text{acc}}/\sqrt{2e}$ is the dimensionless acceleration time normalized by the ion plasma frequency ω_{pi}. This model provides a refinement over previous freely expanding plasma models by taking into account the limited duration of electron injection and sheath field formation. In experiments, the acceleration time is found to scale approximately as $t_{acc} \sim 1.3\tau_{laser}$ (Fuchs et al., 2006).

The resulting ion energy spectrum is approximated by a quasi-thermal distribution with a sharp cutoff at the maximum energy. The number of protons per unit energy is given by Mora (2003) :

Equation 1.31 : TNSA proton yield (isothermal expansion model)

$$\frac{dN}{d\mathcal{E}} = \left[\frac{n_{e0}\,c_s\,t_{\text{acc}}\,S_{\text{sheath}}}{2\,\mathcal{E}\,k_BT_{hot}}\right]^{1/2}\exp\left(-\sqrt{\frac{2\mathcal{E}}{k_BT_{hot}}}\right) \tag{1.31}$$

where $c_s = \sqrt{Z_iT_{hot}/m_i}$ is the ion sound speed. This expression predicts an exponentially decaying spectrum with a sharp cutoff, consistent with experimental observations of proton beams from solid targets irradiated at relativistic intensities.

Building on the theoretical description of the ion spectrum, experimental studies have sought to directly characterize the ultrafast sheath fields that govern the acceleration process. Femtosecond-resolved measurements using Electro-Optic Sampling have revealed that the accelerating sheath field in TNSA can reach magnitudes of up to ~ 0.6 TV/m at the target surface (Pompili et al., 2018). These measurements, performed at the *FLAME* 100 TW laser using a ZnTe crystal, showed that the field amplitude scales with the laser energy approximately as $E_T \propto \mathcal{E}_L^{0.30}$, and that the charge expansion across the target surface occurs at nearly the speed of light. Using the experimentally

established scaling law, we estimate sheath electric fields of approximately 0.7 TV/m for *ALLS* (3.2 J), 2.1 TV/m for *Titan* (130 J), and 2.6 TV/m for *Apollon-LULI* (265 J), confirming the multi-TV/m regime for PW-class systems. These measurements confirm the theoretical predictions of sheath evolution models (Romagnani et al., 2005). Advanced experimental diagnostics, such as proton radiography, have enabled the visualization of sheath dynamics in real-time (Simpson et al., 2021a, 2023). These techniques provide valuable insights into the spatial and temporal evolution of the accelerating fields. Additionally, interferometry-based measurements have been used to characterize the expansion of the plasma sheath and the formation of strong electric fields at the rear surface of thin foil targets (Green et al., 2014a; Feister et al., 2014; Bisesto et al., 2019). The results from these experimental studies have been used to validate theoretical models predicting ion energy distributions and acceleration efficiencies.

1.3.8 Scaling laws in Target Normal Sheath Acceleration

An increase in laser intensity on target generally results in a rise in the maximum proton energy ($\mathcal{E}_{p,max}$) within the TNSA regime, as shown in Figure 1.11. However, the appropriate scaling law governing proton energy as a function of laser irradiance remains a subject of ongoing debate. While Figure 1.11 highlights the influence of laser irradiance, it is clear that other factors such as pulse duration, temporal contrast, polarization, incidence angle, and, equally importantly, target characteristics (including material composition, density, thickness, surface roughness, and geometry) also play critical roles in determining the final proton energy. Several parametric studies have explored how $\mathcal{E}_{p,max}$ depends on laser irradiance, pulse duration, energy, and fluence. Four models (Fuchs et al., 2006; Robson et al., 2007; Flippo et al., 2008; Zeil et al., 2010) for laser-driven proton acceleration offer complementary insights into how the maximum proton energy scales with laser parameters, each developed under specific experimental and theoretical conditions.

Fuchs et al. (2006) introduced a time-limited isothermal fluid model for proton acceleration in the TNSA regime, incorporating ponderomotive scaling to estimate the hot electron temperature and a realistic acceleration time proportional to the laser pulse duration ($t_{\mathrm{acc}} \sim 1.3\tau_{laser}$). This model predicts the maximum proton energy and energy conversion efficiency over a wide range of laser parameters (10^{18}–10^{20} W.cm^{-2}), as confirmed by both experiments and 2D PIC simulations. The model assumes a quasi-one-dimensional, isothermal plasma expansion from the rear target surface and does not account for multi-dimensional effects or temporal electron cooling, yet still provides strong agreement with measured scaling trends for varying laser energy, pulse duration, and target thickness.

In contrast, Robson et al. (2007) extended the experimental investigation of laser-driven proton acceleration up to intensities of 6×10^{20} W.cm^{-2}, revealing that the one-dimensional isothermal plasma-expansion model employed by Fuchs et al. (2006) substantially overestimates the maxi-

mum proton energy in this ultrahigh-intensity regime. Based on their dataset, Robson et al. (2007) established a simple empirical scaling of the form

Equation 1.32 : Empirical proton cutoff scaling (Robson *et al.*)

$$\mathcal{E}_{p,max} = a\, I^{b} \tag{1.32}$$

where $b = 0.5 \pm 0.1$. obtained for aluminum foils of 10–25 μm thickness at a constant pulse duration of 1 ps. The authors further noted only a weak dependence of $\mathcal{E}_{p,\max}$ on the laser pulse duration (1–8 ps) and a systematic increase in proton energy for thinner targets, both attributed to enhanced hot-electron coupling and sheath-field formation. Although no explicit multi-parameter formula was provided, these observations emphasize that a combination of laser intensity, pulse duration, and target geometry influences proton energy scaling. To reconcile experimental results with theory, Robson et al. (2007) introduced a revised two-phase expansion model including time-dependent electron temperature and approximate three-dimensional effects, yielding improved agreement with the experimental data.

Flippo et al. (2008) provided another empirical scaling law from experiments on the *Trident* 200 TW laser, reporting :

Equation 1.33 : Empirical proton cutoff scaling (Flippo *et al.*)

$$\mathcal{E}_{p,max} \propto I^{0.565} \tag{1.33}$$

with observations spanning 10^{18} to 10^{20} W.cm^{-2}. Their results are consistent with the relativistic ponderomotive scaling and support the notion that proton energy is not only a function of hot electron temperature but also of sheath field geometry and temporal coupling.

Zeil et al. (2010) conducted a systematic experimental study using ultrashort laser pulses ($\sim$30 fs) from the 100 TW-class *Draco* Ti : Sapphire laser and demonstrated a linear scaling of maximum proton energy with laser power rather than with intensity. This behavior was attributed to the short acceleration timescale, comparable to the laser pulse duration, and the localized sheath field formation at the rear surface of micron-thick foil targets. The authors provided both analytical modeling and 2D PIC simulations to support this interpretation, emphasizing the role of 3D effects and the importance of maintaining tight focusing and excellent temporal contrast.

Altogether, these models vary in applicability based on pulse duration, intensity regime, target thickness, and temporal contrast, and they highlight the need to account for laser-plasma interaction physics beyond simplistic 1D assumptions as experimental conditions push into new regimes. An alternative class of models assumes that the highest-energy ions behave as test particles accelerated in a quasi-static sheath field, unaffected by their dynamics. These static models depend critically on accurate characterizations of the sheath field, which in turn rely on realistic assump-

tions about the distribution of hot electrons. Passoni et al. (2008) developed a theoretical model describing ion acceleration in the early stage of the TNSA process, focusing on the quasistatic electric field formed by ultrarelativistic bound electrons at the rear surface of thin targets. By considering only electrons with negative total energy (i.e., trapped), the model predicts a finite accelerating sheath and avoids the divergences present in earlier fluid-based expansions. Analytical expressions for the electric field profile, maximum ion energy, and proton spectra are derived and show good agreement with a broad set of experimental data (Passoni et al. (2010) : 0.1–500 J, $I_L \sim 10^{18}$–10^{20} W.cm^{-2}), typically within 10% error. The model also captures enhancements observed with ultrahigh-contrast pulses and provides scaling laws for future systems targeting proton energies up to 300 MeV, relevant to hadron therapy.

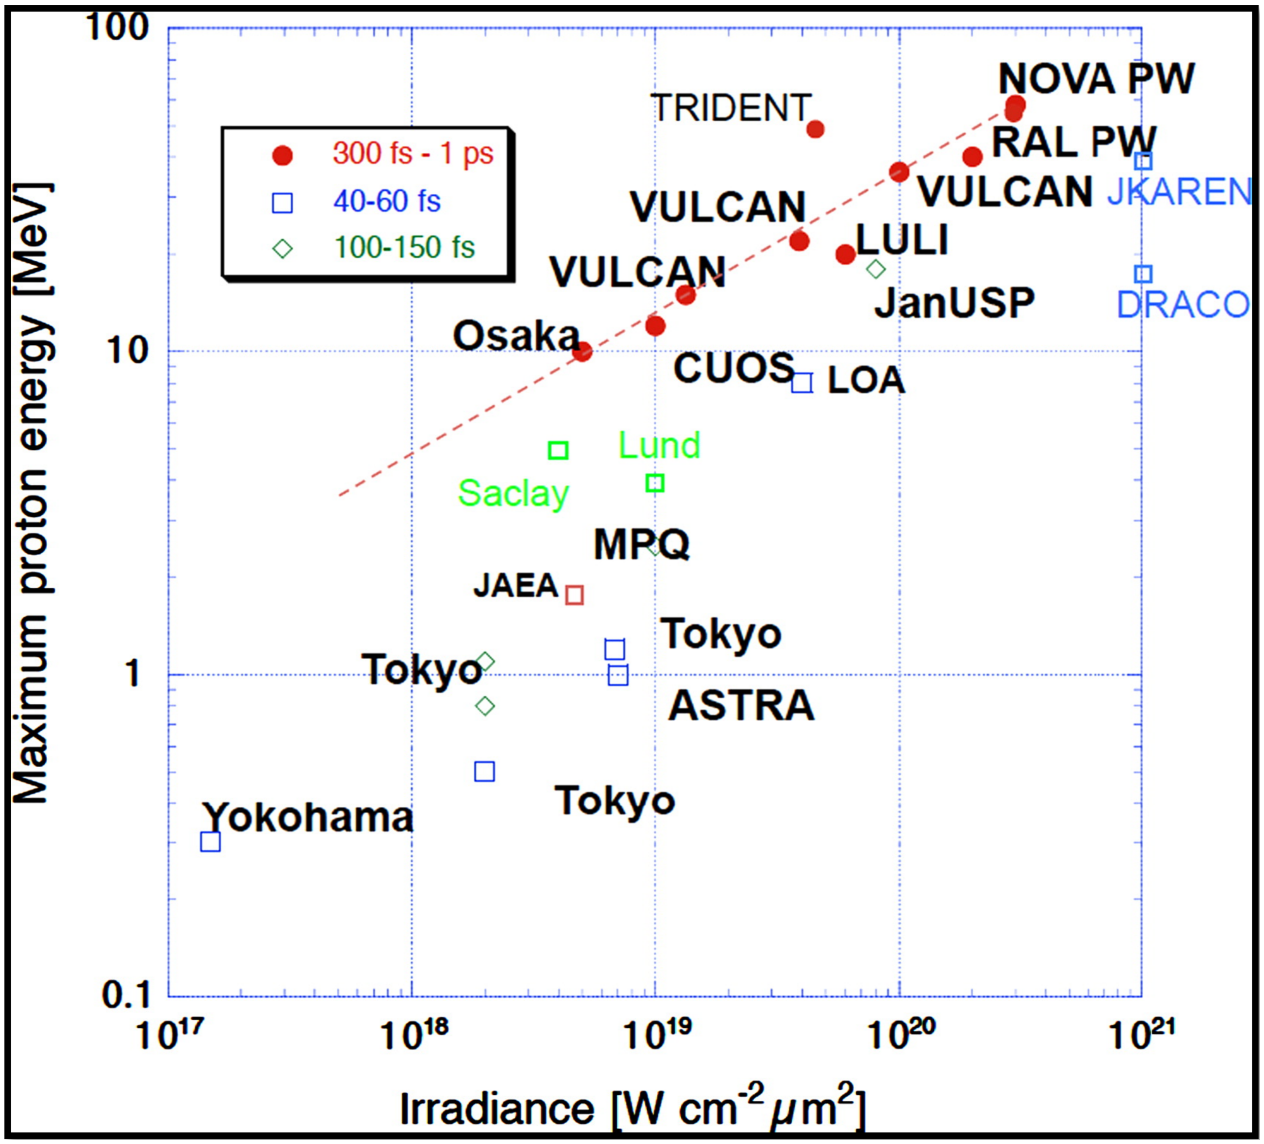

FIGURE 1.11 : Maximum proton energy versus laser irradiance.

Maximum proton energy from laser-irradiated metal foils for experiments on different laser systems as a function of the laser pulse irradiance, grouped in three different ranges of pulse durations. References to the relevant experiments can be found in Borghesi et al. (2006), except for the points labeled Saclay (Ceccotti et al., 2007) and Lund (Neely et al., 2006) marked in light green. This figure has been adapted from Macchi et al. (2013), with the publisher's permission.

1.3.9 Radiation Pressure Acceleration

If the laser intensity climbs above 10^{20} W.cm^{-2}, another acceleration mechanism comes into competition with TNSA when using overdense targets, namely the Radiation Pressure Acceleration (RPA). The RPA mechanism was first suggested in the work of Esirkepov et al. (2004) for intensities in excess of 10^{23} W.cm^{-2}, in an intensity range where it strongly predominates over TNSA.

In this mechanism, the momentum carried by the electromagnetic wave is so high that it pushes directly on the target, providing collective thrust to the particles. This results in a forward-oriented beam of particles emerging from the radiation pressure exerted upon the EM wave's impact with the front target surface. RPA is only effective for very thin targets (typically sub-micrometric in thickness). If the impact is sufficiently strong, the accelerated ions from the front surface in the laser's forward direction can reach velocities exceeding those of ions accelerated at the rear surface through TNSA, while also exhibiting a much lower energy dispersion.

A way to further increase the predominance of RPA is to use circularly polarized laser pulses to suppress hot electron generation — particularly by inhibiting $\vec{J}\times\vec{B}$ heating — and thereby preserve only the drift term in the longitudinal momentum equation (Robinson et al., 2008).

The radiation pressure P_{rad} of a laser pulse of peak intensity I_0 upon impact with a target can be expressed as (Macchi et al., 2013) :

Equation 1.34 : Radiation pressure on a target by a laser pulse

$$P_{\text{rad}} = (1 + \mathcal{R} - \mathcal{T})\frac{I_0}{c} = (2\mathcal{R} + \mathcal{A})\frac{I_0}{c} \tag{1.34}$$

where $\mathcal{R}$, $\mathcal{T}$, and $\mathcal{A}$ are the reflection, transmission, and absorption coefficients, respectively. In equation 1.34, the conservation of energy implies that $\mathcal{R} + \mathcal{T} + \mathcal{A} = 1$. In the ideal case of a "perfect" relativistic mirror ($\mathcal{T} = \mathcal{A} = 0$, $\mathcal{R} = 1$), this leads to a maximum radiation pressure of $P_{\text{rad}}^{\text{max}} = 2I_0/c$. where $\mathcal{R}$, $\mathcal{T}$, and $\mathcal{A}$ are the reflection, transmission, and absorption coefficients, respectively.

It is possible to distinguish two main regimes of RPA : one for thick targets, called the Hole Boring (HB) regime, and another for ultrathin targets, called the Light Sail (LS) regime. Schematic representations of both regimes are provided in Figure 1.12, and further details are given below.

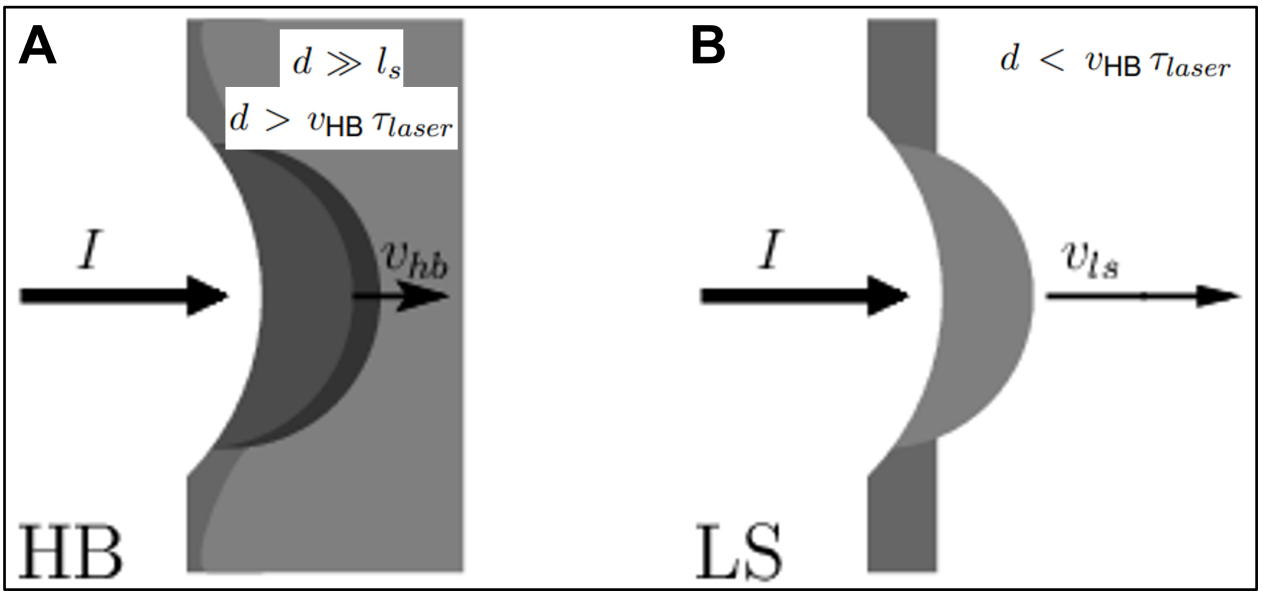

FIGURE 1.12 : Comparison of Hole-Boring and Light-Sail RPA regimes.

Schematic diagrams of the RPA. (A-B) Two distinct acceleration cases within RPA are the Hole Boring (A) and Light-Sail (B) regimes. These figures have been adapted from Macchi (2017), with the author's permission.

Hole Boring regime

The Hole Boring (HB) regime (also called the laser piston regime) occurs when the laser pulse does not penetrate through the target, as a result of using a thick target (i.e., much thicker than the laser penetration skin depth, $d \gg l_s$), and pushes the particles to continuously increase the local density from its pressure. This steepens the density of the target as the laser acts on the surface like a piston. This ultimately bores a hole in the plasma and accelerates particles up to a velocity called the hole boring velocity v_{HB}, thereby defining the target thickness condition $d > v_{\mathrm{HB}}\,\tau_{laser}$, with τ_{laser} being the laser pulse duration. Estimations of the HB velocity are found by balancing the electromagnetic and mass momentum fluxes in planar geometry (Robinson et al., 2009; Schlegel et al., 2009). This leads to the following kinetic energy scaling for ions :

Equation 1.35 : Hole boring (RPA) ion kinetic energy scaling

$$\mathcal{E}_K^{\mathrm{HB}} = 2\, m_e c^2\, Z\, \frac{n_c}{n_e}\, a_0^2 \tag{1.35}$$

Hence, we see from equation 1.35 that the ion kinetic energy scales as $\mathcal{E}_{\mathrm{K}}^{\mathrm{HB}} \propto a_0^2 \propto I_0$, compared to $\mathcal{E}_{\mathrm{K}}^{\mathrm{TNSA}} \propto a_0 \propto \sqrt{I_0}$ for the TNSA mechanism, explaining why RPA attracts significant attention in the laser acceleration community.

It is also worth noting that $\mathcal{E}_{\mathrm{K}}^{\mathrm{HB}} \propto 1/n_e$, and thus high electron densities are not conducive to high ion energies in the HB-RPA regime. The target density must therefore be reduced close to

$n_e \sim n_c$ but not below, since the target would become transparent, thereby reducing the reflection coefficient and the strength of the impact. Moreover, reducing the target thickness also helps to increase ion energies by minimizing collisional energy losses.

Promising theoretical estimates using $I_0 > 10^{22}$ W.cm^{-2} and near-critical targets of micrometric thickness predict proton energies exceeding 100 MeV and quasi-monoenergetic beams (Klimo et al., 2008; Macchi et al., 2010; Robinson et al., 2012).

Light Sail regime

If the target thickness d is reduced such that $d < v_{\text{HB}}\,\tau_{laser}$, with v_{HB} being the hole boring velocity and τ_{laser} the laser pulse duration, the ion front reaches the rear side of the target before the end of the laser pulse. In this situation, the entire target is accelerated as a single unit within the focal volume. This regime is known as the Light Sail (LS) regime, and the target behaves like a thin relativistic mirror propelled by radiation pressure. This configuration minimizes particle energy losses by eliminating the background plasma and continuously accelerates the same ion layer that was initially compressed by the laser.

Ion energy estimations in the LS regime yield the following expression for ion kinetic energy (Qiao et al., 2009; Macchi et al., 2013) :

Equation 1.36 : Light Sail (RPA) ion kinetic energy scaling

$$\mathcal{E}_K^{\text{LS}} = m_p c^2 \frac{F^2}{2(1+F)} \tag{1.36}$$

with

$$F = \frac{2Z}{A}\frac{n_c}{n_e}\frac{c\tau_{\text{laser}}}{d}a_0^2.$$

where F is the dimensionless fluence, Z and A are the ion charge and mass numbers, n_c the critical density, and n_e the target electron density. As opposed to the HB regime, the target thickness d explicitly appears in equation 1.36, leading to the scaling $\mathcal{E}_{\text{K}}^{\text{LS}} \propto 1/d$. Therefore, reducing the target thickness enhances energy gain in the LS regime.

However, the LS regime is eventually limited by the onset of relativistic transparency. As the laser intensity increases, the target can become partially transparent to the incident field, thereby reducing the effective radiation pressure. This effect becomes significant when the relativistically corrected critical density drops below the target electron density. When $d \ll \lambda$, the transparency threshold is governed by the condition (Vshivkov et al., 1998) :

Equation 1.37 : Relativistic transparency threshold in the LS regime

$$a_0 > \xi = \frac{n_e\, d}{n_c\, \lambda} \tag{1.37}$$

which shows that reducing the target thickness also lowers the intensity threshold for transparency.

The LS-RPA optimal point occurs when the target becomes transparent at the end of the laser pulse. Inserting equation 1.37 into equation (1.36) and simplifying leads to :

Equation 1.38 : Optimal ion energy in the LS regime at the transparency threshold

$$\mathcal{E}_K^{\mathrm{LS,opt}} = 2\pi\, m_e c^2\, \frac{Z}{A} \left(\frac{c\, \tau_{laser}}{\lambda}\, a_0^2 \right)^2 \tag{1.38}$$

Current record laser intensities (Yoon et al., 2019) suggest optimal target thicknesses on the order of tens of nanometers, which is now achievable with state-of-the-art target manufacturing techniques. For instance, considering $I_0 = 10^{22}$ W.cm^{-2} (i.e., $a_0 \approx 100$) with a solid target density of $n_e = 400 n_c$ and a wavelength $\lambda = 1\,\mu$m yields an optimal thickness of $d_{\mathrm{opt}} \approx 80$ nm.

Despite these advances, producing and maintaining ultra-thin targets at high temporal contrast remains technically challenging. Indeed, a laser-to-prepulse contrast as high as 10^{10} at peak intensities of 10^{23} W.cm^{-2} corresponds to prepulse or ASE levels around 10^{13} W.cm^{-2}. Such intensities already exceed the ionization threshold of most solid targets and can pre-ionize or destroy the foil before the main pulse arrives, undermining the LS acceleration process.

1.3.10 Relativistically induced transparency-driven acceleration

Relativistically induced transparency (RIT) arises when an ultra-intense laser drives electrons to relativistic velocities, increasing their effective mass and thereby raising the transparency threshold so that an initially overdense plasma with density $n_e \gtrsim n_c$ becomes transmissive once the effective critical density grows to γn_c (Palaniyappan et al., 2012; King et al., 2023; Fernández et al., 2017). In ultrathin, rapidly expanding foils, this enables laser propagation into and through the target, reshaping absorption and field structures that strongly affect ion acceleration - most notably in the so-called Break-Out Afterburner (BOA) regime (Yin et al., 2006, 2007; Henig et al., 2009). In BOA, the laser penetrates the target once transparency is reached, enabling sustained volumetric electron heating and stronger ion acceleration. This regime can produce a denser and hotter sheath compared to TNSA, and benefits from reduced collisional energy losses. Experiments and simulations have shown that the onset and duration of RIT depend sensitively on target thickness and laser contrast, with improved contrast shifting the optimal thickness and modestly impacting the peak proton energy (Dover et al., 2023). Moreover, kinetic effects such as electron heating from

finite pulse rise times raise the effective transparency threshold above cold-fluid predictions (Ferri et al., 2019b), while finite ion motion can further modify the transition dynamics. RIT can also act as a ”relativistic plasma aperture” locally boosting on-axis intensity through diffraction of transmitted light and enabling new regimes of ion acceleration and strong-field physics (Jirka et al., 2021). Together, these studies establish RIT as a transient, controllable pathway linking target engineering, laser contrast, and hot-electron phase-space to enhanced ion and radiation sources (Macchi et al., 2013).

1.3.11 Magnetic vortex acceleration

Magnetic vortex acceleration (MVA) is a mechanism of ion acceleration that emerges when an ultra-intense laser pulse propagates through a near-critical or underdense plasma, generating a strong, self-consistent magnetic vortex structure in the wake of the laser pulse (Nakamura et al., 2010; Park et al., 2019). In this regime, the ponderomotive force of the laser expels electrons radially, forming a positively charged channel that sustains multi-megagauss azimuthal magnetic fields due to the return currents at the channel boundary (Zhidkov et al., 2007). As the laser exits the plasma, these magnetic fields collapse, converting magnetic energy into an intense longitudinal electric field that accelerates ions located near the rear plasma-vacuum interface (Nakamura et al., 2010). The process is highly dependent on plasma density, target thickness, and laser focal geometry. Optimal MVA conditions are typically achieved in near-critical plasmas where the laser can fully penetrate but still drive strong collective electron motion. Particle-in-cell simulations and experiments have shown that MVA can generate quasi-monoenergetic ion beams with energies exceeding hundreds of MeV, offering an alternative to TNSA and RPA in the relativistic transparency regime in RIT-driven acceleration. Recent studies have highlighted the possibility of controlling the vortex topology and associated fields through pre-shaped plasma channels or density gradients, paving the way for compact, high-repetition-rate sources of energetic ions (Garten et al., 2024; Tazes et al., 2024). Overall, MVA links the relativistic laser–plasma dynamics of underdense regimes with efficient ion acceleration through the conversion of magnetic energy into directed electrostatic potential energy.

1.3.12 Collisionless shock acceleration

Collisionless shock acceleration (CSA) is a prominent mechanism for ion acceleration in high-intensity laser–plasma interactions, occurring when a laser pulse drives a sharp electrostatic shock front through an underdense or near-critical plasma (Silva et al., 2004; Haberberger et al., 2012). In this process, the laser’s ponderomotive force and electron pressure gradients launch a nonlinear ion-acoustic wave that steepens into a collisionless shock, mediated by collective electromagnetic fields rather than binary collisions (Forslund et al., 1971; Tidman, 1971). Ions in the upstream

plasma that encounter this moving potential barrier are reflected and accelerated to twice the shock velocity in the laboratory frame, yielding quasi-monoenergetic ion beams with narrow energy spreads and low divergence (Haberberger et al., 2012). The efficiency and energy of CSA strongly depend on the Mach number $M = v_s/c_s$, plasma temperature gradients, and the laser–plasma coupling conditions, which determine the shock velocity v_s and stability (Fiúza et al., 2012). Experiments using CO_2 and Ti : Sapphire lasers have demonstrated CSA-driven proton beams up to tens of MeV with energy spreads below 10%, confirming its potential for controlled, high-quality ion sources (Romagnani et al., 2008; Haberberger et al., 2012). Recent studies have explored how tailored plasma density ramps and multi-pulse irradiation can stabilize shock formation, extend its lifetime, and enhance ion energies (Pak et al., 2018; Kumar et al., 2019; Marquès et al., 2024). Overall, CSA represents a robust, tunable mechanism linking laser-driven plasma hydrodynamics and collective field structures to the production of monoenergetic, high-brightness ion beams.

1.3.13 Limitations and pathways toward stable, high-repetition-rate laser-driven ion acceleration

Despite the rapid progress in LDIA, a major constraint lies in the low repetition rates of PW-class laser systems, which typically operate at or below 10 Hz (see Figure 1.4) due to thermal management challenges, damage thresholds of optical components, and the need for target replacement between shots. This restricts the average particle flux and hinders applications requiring high-dose delivery or statistically significant datasets. In addition, shot-to-shot reproducibility remains a significant challenge : slight variations in laser pulse quality, preplasma formation, or target microstructure can lead to substantial fluctuations in beam energy, direction, and charge (yield). This inherent instability complicates data interpretation and reduces the reliability as well as the overall implementation of laser-driven ion sources in real-world settings. Current research efforts are addressing these issues through innovations in laser technology (e.g., OPCPA systems), advanced target engineering (e.g., structured and layered foils), and real-time feedback optimization (e.g, wavefront sensors, laser chirp controllers...) to improve stability, spectral control, and system repetition rates. This is why now strong interest is given to laser-driven sources of smaller scale (i.e., generating lower proton energies) but providing higher reliability and repetition rate for applications, able to run at a few Hz to tens of Hz as typically found with commercial TW-scale Ti : Sapphire lasers. Consequently, optimizing laser–target interactions for reliable and high-quality ion beams requires not only the fine-tuning of laser parameters (energy per pulse, intensity, pulse duration, contrast, polarization, and incidence angle) but also an equally careful selection of the target properties, including material composition, density, thickness, surface roughness, and macroscopic geometry.

1.3.14 Repetition rate constraints and target technologies in laser-driven ion acceleration

In TNSA, the properties of the target - specifically its material composition, thickness, and density - play a central role in determining the efficiency and quality of the accelerated ion beams. Among these, target thickness has been the subject of numerous studies, which have explored its impact on hot electron transport and sheath field formation (Dong et al. (2003); Esirkepov et al. (2006); Zeil et al. (2010), among others). However, achieving optimal performance in laser-driven ion acceleration also depends on overcoming several experimental constraints. One major challenge stems from the stringent alignment requirements imposed by the use of large numerical apertures, which are essential for focusing the laser to ultra-high intensities. Additionally, the use of solid targets presents practical limitations, particularly regarding the need to refresh the target after each laser shot in high-repetition-rate setups.

The choice of target material affects various factors, including electron density, ionization potential, and the resultant electric fields, all of which play pivotal roles in the acceleration process. Targets composed of high atomic number (high-Z) materials, such as gold (Au) or tantalum (Ta), exhibit higher electron densities and ionization potentials compared to low-Z materials like carbon or plastic. The increased electron density in high-Z materials enhances the production of hot electrons when irradiated by high-intensity laser pulses. More electrons are establishing intense sheath electric fields at the rear surface, which are crucial for ion acceleration. However, high-Z materials can also introduce complexities such as increased bremsstrahlung radiation, which may affect the overall acceleration efficiency (Nishiuchi et al., 2020; Măgureanu et al., 2022). Incorporating high-Z elements into low-Z matrices, such as embedding gold nanoparticles into polymer substrates, has been explored to optimize TNSA performance. This approach aims to combine the favorable properties of both materials, enhancing electron density while maintaining manageable ionization thresholds. Experimental studies have demonstrated that such composite targets can lead to increased hot electron generation and, consequently, stronger sheath fields, resulting in more efficient ion acceleration (Torrisi et al., 2017, 2019). The presence of surface contaminants, often hydrocarbons or water vapor, can significantly impact the ion species and energies in TNSA. These contaminants, typically low-Z materials, are ionized and accelerated preferentially due to their lower ionization potentials. Implementing techniques such as fs-laser desorption to remove these contaminants has been shown to modify the ion acceleration dynamics, reducing proton acceleration and enhancing the acceleration of target material ions (Hoffmeister et al., 2013). Utilizing nanostructured targets, such as those with nanocone or nanohole arrays, has been shown to enhance laser absorption and hot electron production. These structures increase the effective interaction area and create localized field enhancements, leading to more efficient TNSA processes. Simulations have indicated that such nanostructuring can result in higher proton cutoff energies compared to flat targets. Here are a few studies reporting the use of nanostructured targets for TNSA-enhanced process : Barberio et al. (2016); Vallières et al. (2017); Thiele et al. (2019); Val-

lières et al. (2021). The material composition of the target is a critical parameter in TNSA, influencing electron dynamics, sheath field formation, and ion acceleration efficiency. Careful selection and engineering of target materials, including the use of high-Z elements, composites, and nanostructuring, can significantly enhance the performance of laser-driven ion acceleration systems.

Experimental studies have demonstrated a strong correlation between target thickness and the maximum energy of accelerated protons in the TNSA regime. While thinner targets often enable higher proton energies by reducing the scattering and energy loss of hot electrons as they transit the target, this trend holds only within a specific range. For example, a 0.65-µm-thick aluminum target produced protons with a maximum energy of 3.5 MeV, compared to only 1.25 MeV for a 12.5-µm-thick target (Su et al., 2013). It is important to distinguish between nanometric and micrometric targets, as their behavior differs markedly depending on laser contrast and pulse duration. There exists an optimal target thickness below which the acceleration efficiency decreases. However, for ultrathin targets (a few tens of nanometers), a high temporal contrast is essential to avoid premature expansion or destruction of the target by the laser pedestal or prepulse, which would suppress efficient sheath formation. In such cases, insufficient hot electron generation can occur, or the rear side of the target may deform or expand before the main pulse arrives, leading to inefficient ion acceleration. Moreover, ultrathin targets can become relativistically transparent or even induce a transition toward different acceleration regimes such as RPA, especially under conditions of high contrast and circular polarization (Macchi et al., 2013). In RPA, the laser pulse directly transfers momentum to the entire target, which can result in quasi-monoenergetic ion spectra and higher acceleration efficiencies, but this regime is highly sensitive to laser and target parameters. Additionally, target thickness influences beam divergence. Thicker targets tend to yield more collimated beams due to their resistance to deformation by pre-plasma-induced shock waves, whereas thinner targets often generate broader angular distributions (Lundh et al., 2007; Esirkepov et al., 2014; Ivanov et al., 2014). This sensitivity is again linked to contrast conditions and the presence of prepulses, which can pre-ionize or expand the target rear surface. Finally, a critical thickness threshold exists, below which targets become transparent or fail to support adequate charge separation. In experiments with high-contrast pulses, this threshold has been well characterized (Lécz et al., 2022), reinforcing the idea that optimal performance is obtained only when both the laser contrast and the target thickness are jointly considered.

The density of the target material also significantly influences the acceleration process. Lower-density target materials can achieve higher maximum proton energies. This enhancement is attributed to the more efficient generation and transport of hot electrons in low-density materials, leading to stronger sheath electric fields at the target's rear surface. However, lower-density targets are more susceptible to deformation due to laser-induced shock waves, which can increase the divergence of the proton beam. Utilizing near-critical density (NCD) targets, where the electron density is close to the critical density for the laser wavelength, has been shown to enhance laser absorption and hot electron production. This enhancement leads to more efficient proton acceleration. Simulations indicate that coupling a solid target with a low-density layer can significantly

increase both the energy and number of accelerated protons (Sgattoni et al., 2012; Tong et al., 2018). Recent experiments using foam targets with densities between those of gases and solids have provided insights into ion acceleration mechanisms. These studies have demonstrated that foam targets can influence the charge states of accelerated ions, affecting their energy deposition and transport properties (Passoni et al., 2014; Ma et al., 2023).

Recent advancements in target design for laser-driven ion acceleration have explored several innovative configurations, including liquid, tape, and cryogenic targets, each offering distinct advantages and challenges. Liquid targets, such as liquid leaf targets, have emerged as promising candidates for high-repetition-rate experiments. Their continuously refreshing surface ensures consistent interaction conditions, which is essential for stable ion acceleration at high shot rates (Schnürer et al., 2005; Morrison et al., 2018; George et al., 2019; Treffert et al., 2022b; Puyuelo-Valdes et al., 2022). Moreover, recent studies have employed artificial neural networks trained on PIC simulation data to develop surrogate models for liquid leaf target experiments, enabling optimization of experimental parameters for maximum ion energy and laser-to-ion conversion efficiency (Schmitz et al., 2023).

Cryogenic targets (Morrison et al., 2012; Gauthier et al., 2016; Margarone et al., 2016; Scott et al., 2018; Polz et al., 2019) have also gained attention, particularly those employing isotopically layered structures. A notable design includes a nanometer-thick layer of pure deuterium at the target-vacuum interface, which facilitates a dual-species plasma expansion (Scott et al., 2018). This configuration enables the generation of composite ion beams with tunable density ratios and allows for the suppression of high-energy protons in favor of spectrally peaked deuteron beams, opening new possibilities for tailored ion beam production.

Finally, high-repetition-rate tape target delivery systems have been developed to meet the demands of ultra-intense laser-matter interactions. These systems employ spooled solid tapes to deliver fresh material to the laser focus for each shot, allowing continuous operation under high-vacuum conditions (Fujii et al., 2003; Nayuki et al., 2003; Noaman ul Haq et al., 2017; Dover et al., 2020b). Recent work by Ehret et al. (2023) demonstrated excellent shot-to-shot stability, confirming that tape targets are a reliable and scalable solution for high-repetition-rate laser applications.

In conclusion, careful selection and optimization of target material composition, thickness, and density are essential for enhancing the efficiency and quality of LDIA. Understanding the interplay of these parameters allows for the design of targets that tailor ion beams, thereby advancing the development of compact and efficient laser-driven ion accelerators.

Figure 1.13 shows an overview of some target delivery systems for high-repetition-rate LDIA experiments.

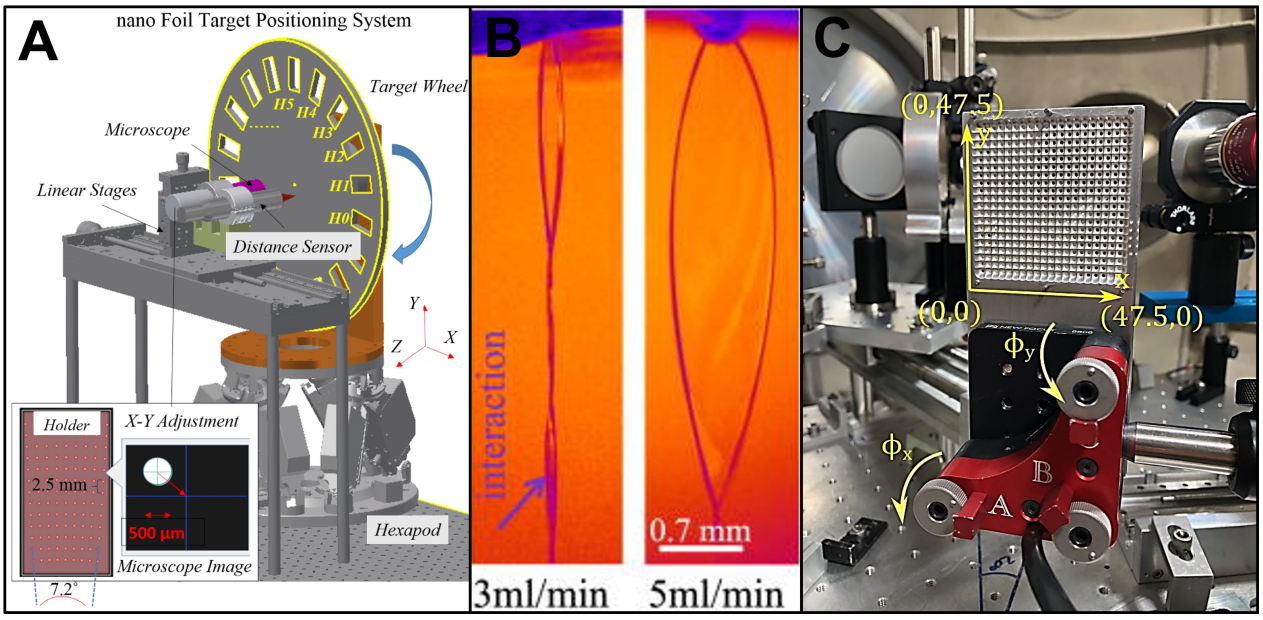

FIGURE 1.13 : Target delivery systems for high-repetition experiments.

Overview of target delivery systems for high-repetition-rate laser-driven ion acceleration experiments. (A) Tape drives (Gao et al., 2017). (B) Liquid jets provide stable, self-refreshing targets suitable for high repetition rates (Treffert et al., 2022a). (C) Multi-target holders allow precise positioning of dense target arrays for efficient experimental throughput (400 shots per pumping cycle) (Catrix et al., 2023). These figures have been adapted with the publishers' permissions.

1.3.15 Diagnostics in laser-driven ion acceleration

In LDIA, the adage "a beamline is only as good as its diagnostics" holds particularly true. Accurate and comprehensive diagnostics are essential not only for unraveling the underlying acceleration mechanisms but also for optimizing performance and tailoring the ion beam properties to meet the demands of specific applications. These diagnostics encompass a range of techniques designed to measure key beam parameters such as energy spectrum, spatial distribution, temporal profile, and charge state. The energy distribution of accelerated ions is typically characterized using Thomson parabola spectrometers (TPS) (see Figure 1.14.A and B.). These devices separate ions based on their charge-to-mass ratio and energy, providing detailed insights into the ion species present and their respective energies (Sadowski et al., 2009; Gwynne et al., 2014; Tata et al., 2017). The TPS, which separates ions via simultaneous deflection in parallel electric and magnetic fields, is particularly well-suited for its near-unambiguous mapping of ion trajectories onto the detector plane. To accommodate a wide range of ion energies and charge-to-mass ratios, both the electric and magnetic field strengths must be adjustable. Adjusting the electric field is relatively straightforward, typically achieved by varying the applied voltage. Several groups have further optimized geometries to enhance the dynamic range of the electric field (Sakabe et al., 1980; Morrison et al., 2011; Alejo et al., 2016; Nedbailo et al., 2023). In contrast, tuning the magnetic field is more complex, especially when relying on permanent magnets, which are commonly used to establish the di-

pole field. Innovations such as dynamic yokes have been introduced to allow adjustable magnetic field strengths (Kojima et al., 2020), and some designs employ electromagnets to adapt the field strength to the analyzed ion species and energy (Harres et al., 2008; Schillaci et al., 2014). These approaches are generally tailored to ions within the MeV energy range. Building on this, Riedlinger et al. (2024) presented an electromagnetic TPS that leverages a fully tunable electromagnet, granting independent control over both deflection axes. This capability enables the spectrometer to capture wide ion spectra with minimized measurement uncertainty. Crucially, the system is designed for vacuum compatibility without requiring active coil cooling. Its compact form factor allows easy integration into experimental setups, preserving space for essential components such as focusing optics, beam dumps, and additional diagnostics.

Accurate measurement of the total charge in ion beams is vital, especially for applications like radiotherapy where dose precision is critical. Integrating Current Transformers (ICTs) offers a noninvasive means to measure the charge of accelerated ion bunches, enabling online dose measurements and facilitating beam tuning. The implementation of ICTs in PW LDIA experiments has demonstrated its effectiveness in real-time charge diagnostics. Geulig et al. (2022) presents the first successful implementation of an ICT as a noninvasive, online charge diagnostic in LDIA experiments at the *BELLA* PW facility. The ICT is used to measure the total charge of ion bunches, enabling real-time dose estimation during radiobiological cell irradiation.

In summary, the array of diagnostic tools for LDIA is both extensive and rapidly advancing. Beyond the TPS, a variety of detectors are employed and combined to capture the full complexity of charged particle emissions. These include Radiochromic Films (RCF), CR-39 track detectors, Imaging Plates (IP), MicroChannel Plates (MCP), and TOF delay lines paired with plastic scintillators, silicon nitride (SiN), or diamond detectors. For example, TOF detectors are employed to measure the temporal profile and energy of ion beams by recording the time taken for ions to reach the detector. This method has been employed to diagnose high-energy laser-driven protons, enabling real-time measurements of key beam parameters (Scuderi et al., 2020; Vallières et al., 2020). Additional techniques involve a compact scintillating fiber imaging spectrometer based on the tomographic reconstruction of proton energy deposition in a layered fiber array (Patel et al., 2024) and pixelated semiconductor detectors (Reinhardt et al., 2011; Würl, 2018). To further enhance the interpretation and efficiency of these diagnostic techniques, some recent efforts have turned toward data-driven approaches. In particular, the integration of Machine Learning algorithms into diagnostic systems represents a significant advancement. For example, neural network-based synthetic diagnostics have been developed to predict the energy spectra of laser-driven protons. By combining variational autoencoders with feed-forward networks, these models can predict energy spectra based on secondary diagnostics of laser-plasma interactions, achieving notable accuracy with limited training data (McQueen et al., 2025).

Each of these diagnostics plays a critical role in unraveling the intricate physics of laser-plasma interactions and in refining ion beam parameters to meet the demanding specifications of both fundamental research and emerging technological applications.

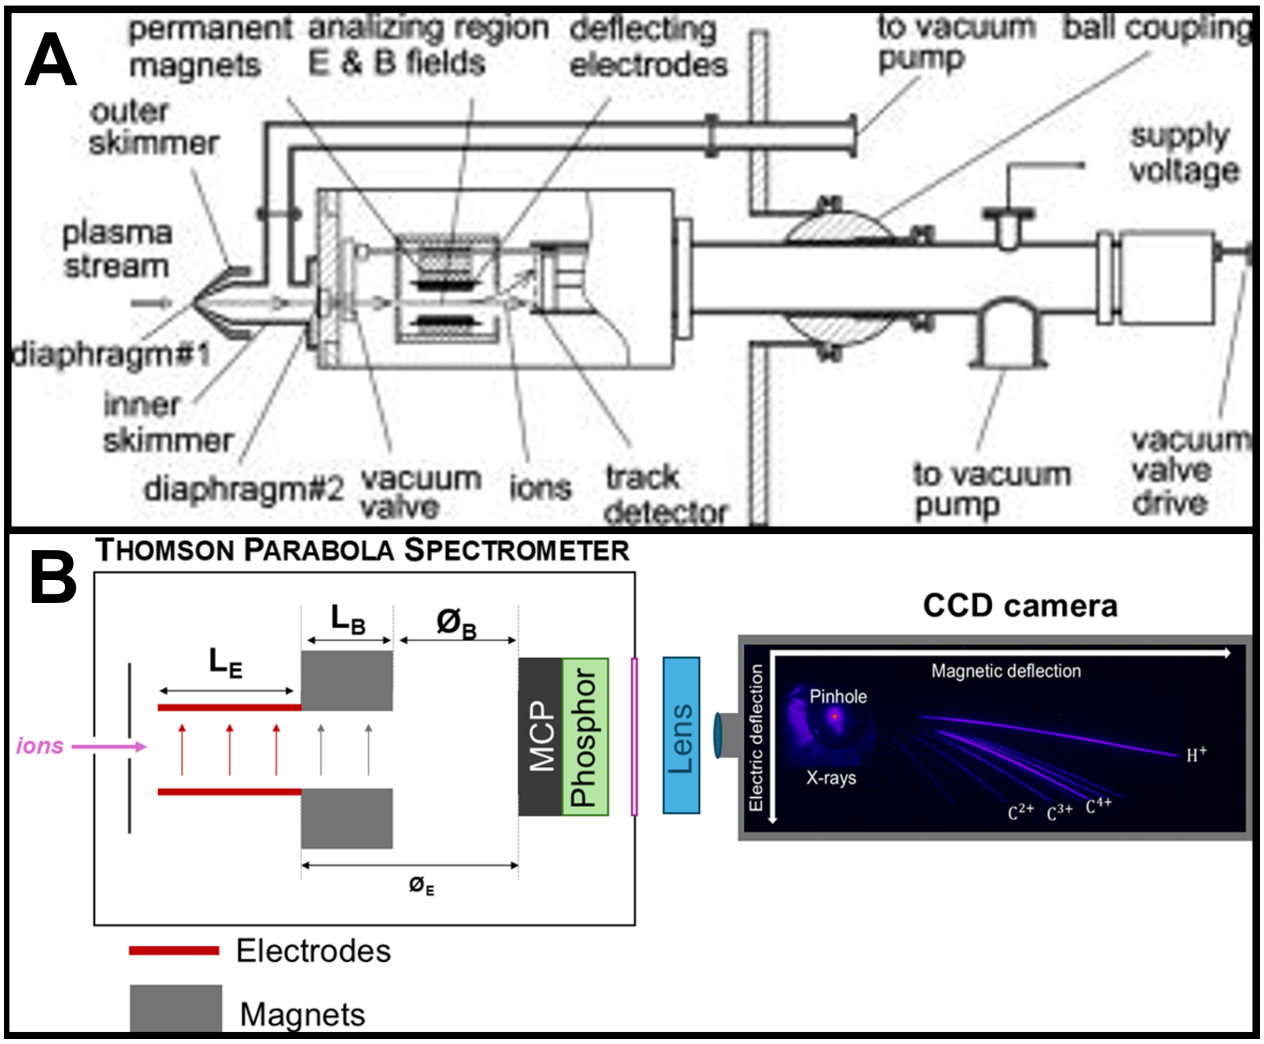

FIGURE 1.14 : Compact Thomson parabola ion diagnostics.

(A) Schematic drawing of the miniature Thomson analyzer (not to scale). This figure has been adapted from Sadowski et al. (2009), with the publisher's permission. (B) Scheme of the TPS at the *ALLS* 150 TW ion beamline.

1.3.16 Applications of laser-driven ion beams

The rapid development of LDIA has been motivated not only by its fundamental interest in high-field science but also by its promising potential across a range of impactful applications. In materials science, laser-driven ions offer unique opportunities to induce structural transitions, phase changes, and defect formation on ultrafast timescales. They are also valuable for highly sensi-

tive elemental analysis, for instance through Particle-Induced X-ray Emission (Sciscio et al., 2017; Barberio et al., 2017b, 2018a, 2019; Puyuelo-Valdes et al., 2021; Boivin et al., 2022), as illustrated in Figure 1.15.A–C. The short duration and high intensity of the LDIA beam enable studies of material behavior under extreme energy densities and strain rates. One of the most prominent and long-standing applications is in inertial confinement fusion (ICF), particularly the fast ignition approach, where a short and intense proton beam is used to rapidly heat a pre-compressed deuterium-tritium (D-T) fuel core to ignition temperatures. In a pioneering study, Roth et al. (2001) proposed that the sheath-accelerated proton beams generated by ultra-intense laser pulses could deliver energy precisely and rapidly into the fuel core, offering a simpler and more localized heating method than relativistic electron transport, which suffers from beam filamentation and divergence issues in dense plasmas (Figure 1.15.D). Experimental campaigns since then have demonstrated proton-induced heating of fusion-relevant materials and validated key aspects of this scheme (Patel et al., 2003; Roth et al., 2005). In hadron therapy, LDIA has attracted attention as a compact and cost-effective alternative to conventional cyclotron and synchrotron-based proton accelerators. Traditional systems are large and expensive, with significant infrastructure and maintenance requirements. Laser-driven protons, by contrast, can be generated using table-top systems and ultra-intense femtosecond lasers, offering the prospect of hospital-scale accelerators (Bulanov et al., 2002; Malka et al., 2004). The ultrashort pulse duration, small transverse emittance, and high particle flux make LDIA beams especially promising for FLASH therapy, which delivers ultrahigh dose rates and has shown normal tissue-sparing effects (Vozenin et al., 2022; Flacco et al., 2024). Although challenges remain in terms of beam stability, energy spread, and repeatability, ongoing advances in beam transport optics and target engineering continue to narrow the gap toward clinical implementation. Another important area is proton radiography, where laser-driven protons act as probing particles for dynamic field structures and dense matter. Their sensitivity to transient electric and magnetic fields, combined with high penetration and ultrafast duration, make them ideal for imaging plasmas, shock fronts, and current filaments with high spatial and temporal resolution. Borghesi et al. (2002) and Mackinnon et al. (2004) demonstrated the use of TNSA-driven protons for diagnosing electrostatic fields in laser-target interactions, revealing filamentation, plasma expansion dynamics, and return currents. This diagnostic capability has been extensively used in high-energy-density physics (HEDP) and laboratory astrophysics (Romagnani et al., 2005; Sarri et al., 2012). LDIA is also a promising driver for compact radioisotope production and applications in nuclear medicine. Using proton beams to induce nuclear reactions in thin targets, short-lived isotopes relevant for positron emission tomography (PET) can be produced on demand. For example, Ledingham et al. (2004) demonstrated for the first time the production of clinically relevant positron-emitting isotopes (^{11}C and ^{18}F) using a high-power PW laser, highlighting a potential alternative to conventional cyclotrons for PET imaging. It also reports the successful synthesis of 2-[^{18}F]FDG, a radiopharmaceutical (specifically a radiotracer), from laser-generated ^{18}F, and discusses the future viability of compact, high-repetition-rate laser systems for on-site isotope generation, which could significantly reduce the infrastructure and cost barriers in nuclear medicine (Figure 1.15.E).

Altogether, these applications - spanning fusion energy, medical therapy, diagnostics, radioisotope production, and materials research - are major drivers of LDIA research. Continued developments in laser technology, target design, and beam control are expected to further accelerate progress toward practical deployment of LDIA systems.

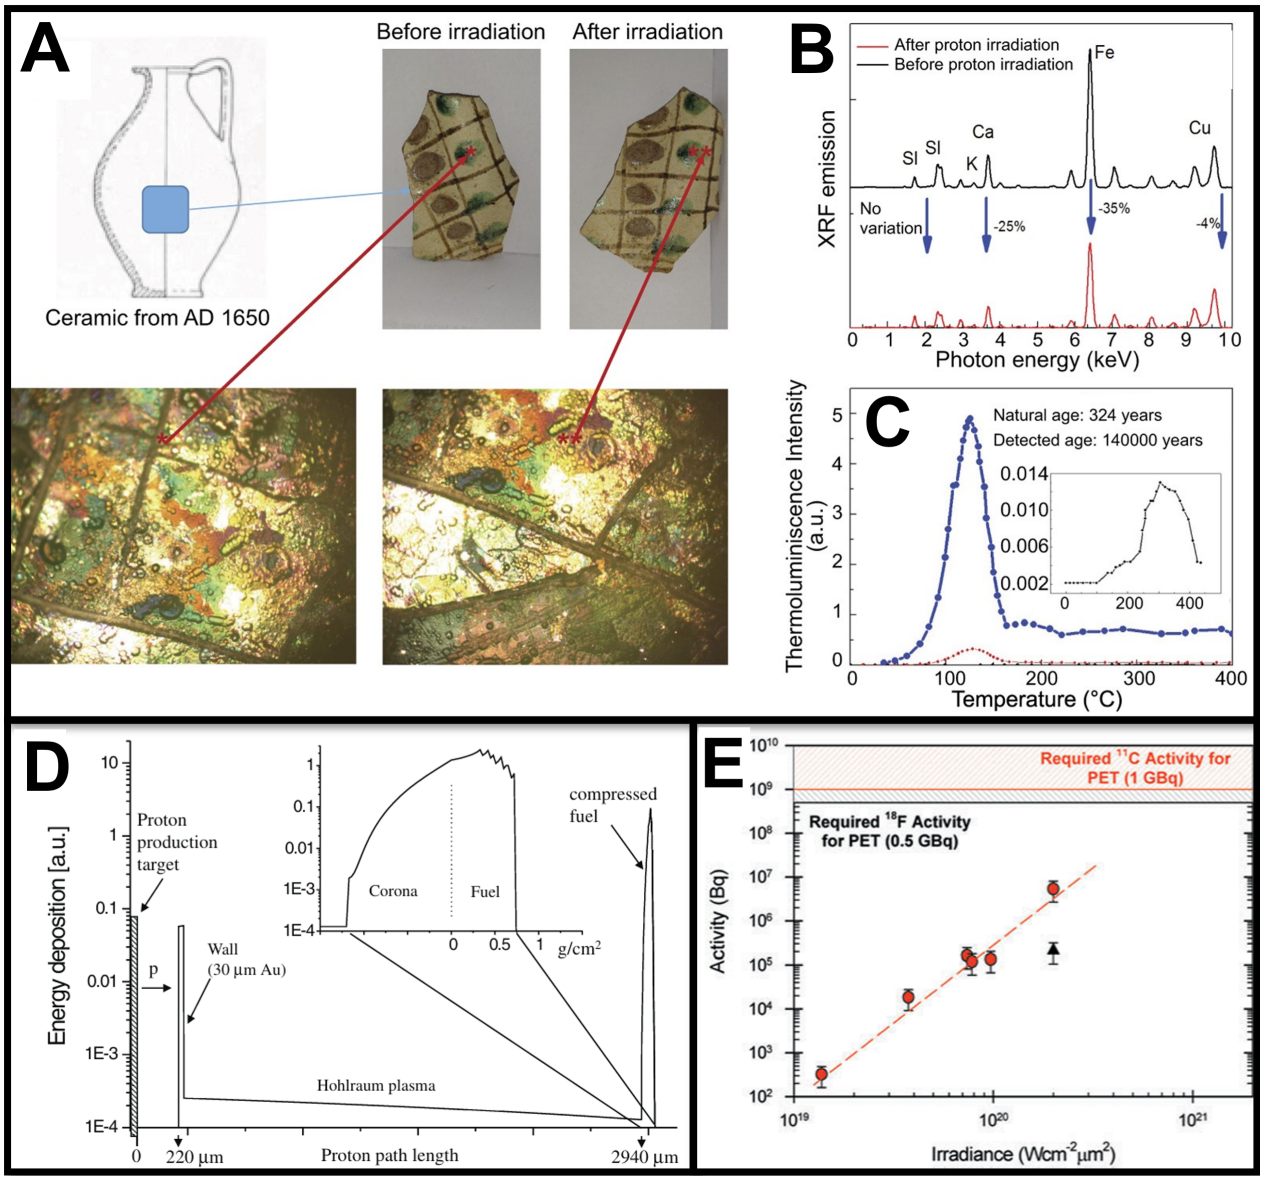

FIGURE 1.15 : Applications of laser-driven ion beams.

(A) Archaeological situ and details about the ceramics used for testing the damaging effect of the laser-generated protons during the one-shot PIXE; Images show the artifact before and after the irradiation; (B) XRF optical image of the sample before and after proton irradiation; (C) Results of the dating process performed before and after irradiation. (A), (B) and (C) have been adapted from Barberio et al. (2017b), with the publisher's permission. (D) Energy deposition of protons between 15 and 23 MeV in a NIF-type indirect drive target. (D) has been adapted from Roth et al. (2001), with the publisher's permission. (E) Total activity (front + back) generated by a single laser shot for both ^{11}C and ^{18}F as a function of laser irradiance with pulse energies from 15 to 300J. The circles refer to ^{11}C production, and the single triangular point for ^{18}F was measured at the highest energies. The hatched areas at the top of the graph indicate the level of required ^{18}F activity from which an ^{18}F-FDG patient dose would be generated, the required ^{11}C activity, e.g. in the form of [^{11}C]CO. (E) has been adapted from Ledingham et al. (2004), with the publisher's permission.

1.4 Conclusion

This chapter has laid the foundational groundwork for understanding laser-driven particle acceleration by examining the interplay between high-power laser technologies and electromagnetic theory. The development of CPA has enabled the generation of ultrashort, ultra-intense laser pulses, revolutionizing our ability to explore extreme field interactions and compact particle acceleration. These advancements have led to the emergence of laser–plasma acceleration schemes that achieve electric field gradients several orders of magnitude higher than those of conventional RF-based accelerators.

Focusing primarily on the TNSA mechanism, we explored how laser-generated hot electrons establish intense sheath fields capable of accelerating ions to multi-MeV energies. The theoretical underpinnings of TNSA, including Maxwell's equations, plasma expansion dynamics, and energy scaling laws, were presented alongside experimental observations and diagnostic techniques. We also briefly introduced alternative regimes such as RPA, CSA, MVA and BOA, which become increasingly relevant at higher intensities or with ultrathin targets.

Moreover, practical considerations such as laser contrast, target design, repetition rate limitations, and diagnostic requirements were discussed in the context of building stable, high-performance ion sources. The chapter highlighted that the performance of laser–plasma accelerators is not dictated solely by laser intensity, but rather by a delicate balance of laser parameters, target properties, and control over electromagnetic field structures.

Collectively, these insights establish a coherent framework for understanding and optimizing the core mechanisms of LDIA and provide the theoretical and technological basis for the experimental strategies, machine-learning methods, and optimization techniques developed in the chapters that follow.

2 MACHINE LEARNING IN LASER-DRIVEN PARTICLE ACCELERATION

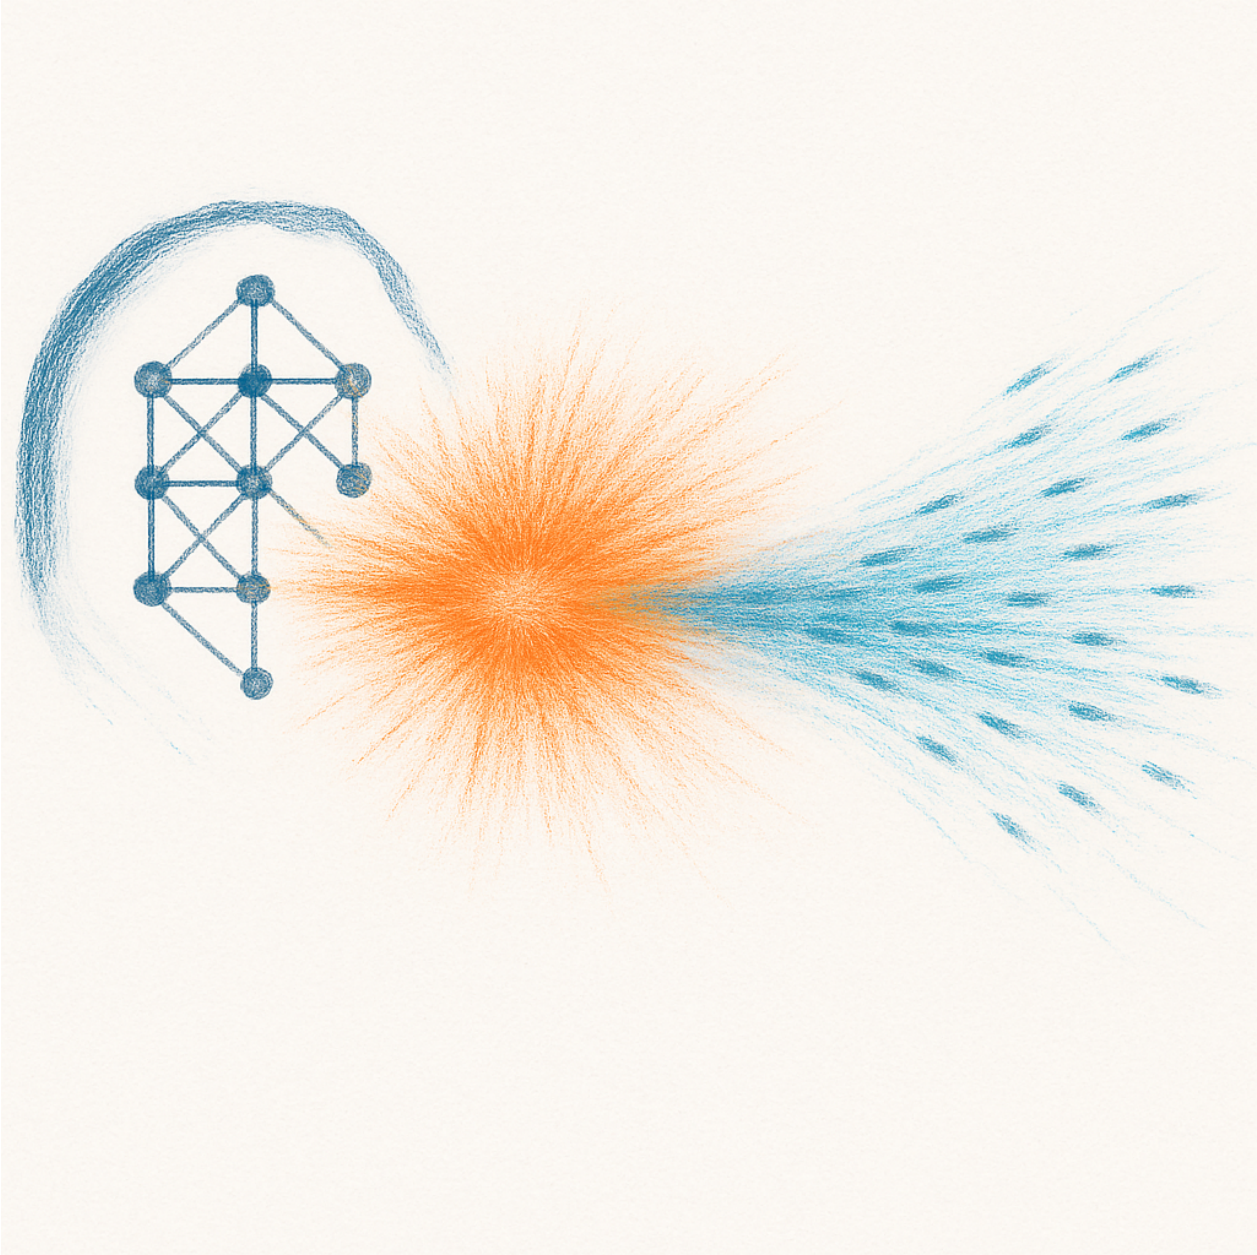

Data-driven control guiding laser–plasma acceleration. Original image created by the author (Elias Catrix). No external permission is required; use authorized by the author.

Foreword — Despite substantial advances in the field, the nonlinear, multi-parameter nature of laser–plasma interactions continues to present significant challenges for both physical understanding and experimental control.

To respond to these challenges, the rise of Machine Learning (ML) driven by advances in computing, algorithms, and data availability has introduced powerful tools for modeling complex systems, inferring structure from limited data, and optimizing experiments where first-principles approaches fall short. These capabilities are particularly well-suited to laser-plasma acceleration, where experimental constraints, limited diagnostics, and multi-parametric control demand adaptive strategies.

This chapter introduces ML concepts and examines its emerging role in plasma physics with a focus on laser-driven ion acceleration. Recent studies are reviewed to illustrate how supervised learning, especially Gaussian process regression, Bayesian optimization, Random Forests and neural networks, have been applied to improve predictive modeling, enhance beam stability, and guide experimental optimization.

As laser systems advance toward higher repetition rates and increasing operational complexity, ML is poised to play a central role in enabling next-generation accelerators to operate with greater efficiency, adaptability, and autonomy. This chapter provides both the conceptual background and the practical context for integrating ML methodologies into high-field plasma physics.

2.1 What is Machine Learning ?

Machine Learning (ML) is a subfield of Artificial Intelligence (AI) (see Figure 2.1) concerned with developing algorithms that allow computers to learn patterns from data and make decisions or predictions without being explicitly programmed for each specific task. As defined by Tom Mitchell in a foundational text (Mitchell, 1997) : "A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E". Unlike traditional programming, where rules are hand-crafted by humans, ML systems construct their internal logic based on data, using statistical inference to capture complex patterns. This paradigm has become especially powerful in areas where analytical models are difficult to derive or where the volume and complexity of data exceed human interpretability.

ML is commonly divided into three broad types (Bishop et al., 2006; Murphy, 2012) :

— **Supervised Learning**
In Supervised Learning, the algorithm is trained on a dataset where each input is associated with a corresponding output (label). The goal is to learn a function that maps inputs to outputs with minimal error. Common examples include classification (e.g., identifying whether an image contains a cat or dog) and regression (e.g., predicting a continuous value such

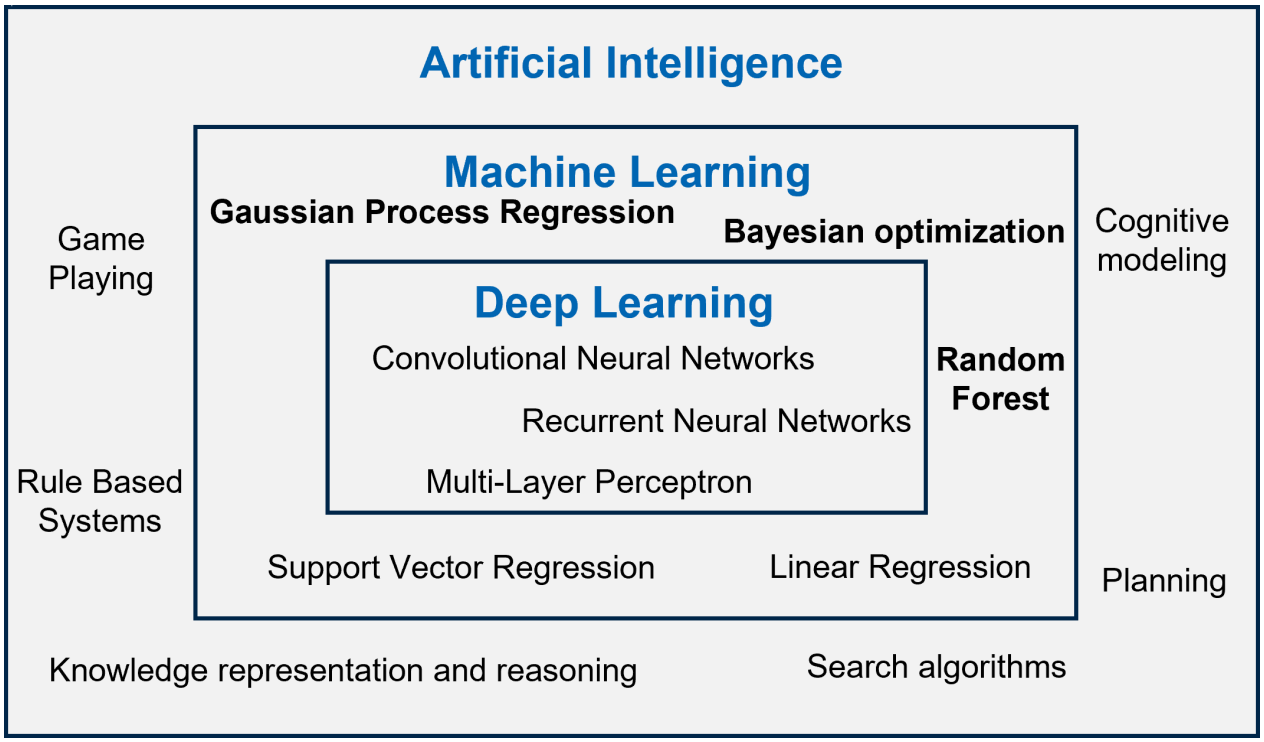

FIGURE 2.1 : Hierarchy of AI, machine learning, and deep learning

Hierarchical relationship between Artificial Intelligence, Machine Learning, and Deep Learning, along with key algorithmic techniques and some used in this doctoral work marked in bold.

as temperature). Techniques such as linear regression, support vector machines, and neural networks are widely used in this setting.

— **Unsupervised Learning**
In contrast, Unsupervised Learning involves datasets without labels. The algorithm must infer structure from the data itself, often by clustering similar examples or reducing the dimensionality of the dataset. Examples include principal component analysis (PCA) and k-means clustering. These methods are especially useful in exploratory data analysis and feature extraction.

— **Reinforcement Learning**
Reinforcement learning (RL) differs fundamentally from the two previous categories. An agent interacts with an environment, receiving feedback in the form of rewards or penalties. Over time, it learns a policy that maximizes cumulative reward. RL has been successfully applied in robotics, resource management, and games (e.g., *AlphaGo* by *Google DeepMind* depicted in Silver et al. (2016)).

A comparison between these three concepts is illustrated in Figure 2.2.

At the heart of most ML methods is a model : a mathematical representation that captures the relationship between inputs and outputs. The process of training such a model typically involves :

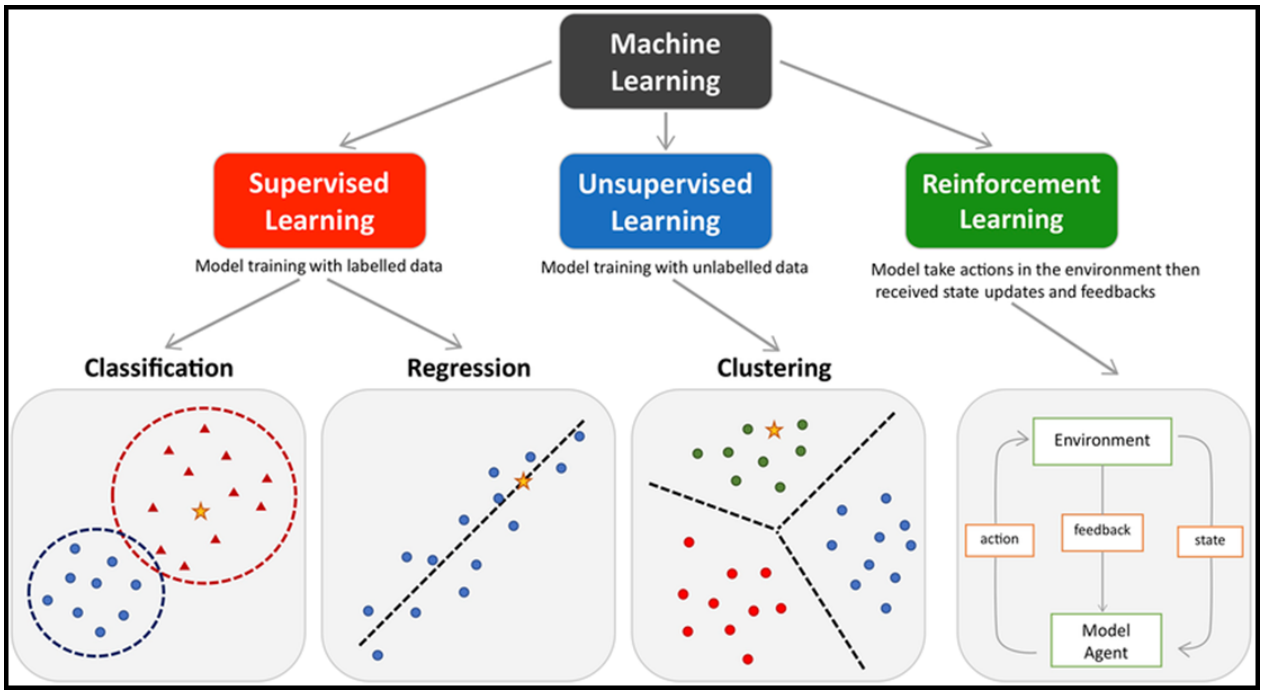

FIGURE 2.2 : Main categories of Machine Learning approaches.

Main approaches include classification and regression under the Supervised Learning and clustering under the Unsupervised Learning. Reinforcement Learning enhances the model performance by interacting with environment. Coloured dots and triangles represent the training data. Yellow stars represent the new data which can be predicted by the trained model. This figure has been adapted from Peng et al. (2021), with the publisher's permission.

1. **Defining a loss function** that quantifies how far predictions are from true labels;
2. **Using optimization algorithms** (e.g., stochastic gradient descent) to minimize this loss;
3. **Evaluating the model** on unseen data to test its generalizability.

A crucial element of modern ML, which workflow is illustrated in Figure 2.3, is the use of neural networks, especially deep learning architectures, which have achieved remarkable success in tasks such as image recognition, natural language processing, and scientific data modeling (LeCun et al., 2015).

Despite its promise, ML has limitations. These include :

— **Data dependence :** Large, high-quality datasets are often required.

— **Overfitting :** Models can become too specialized to the training data and fail to generalize.

— **Interpretability :** Especially in deep learning, models often act as "black boxes."

— **Bias and fairness :** ML models can inherit or amplify biases present in training data.

Addressing these issues is a central focus of current research in the field, and growing interest in explainable AI (Xu et al., 2019) aims to enhance transparency and trust in ML-driven systems.

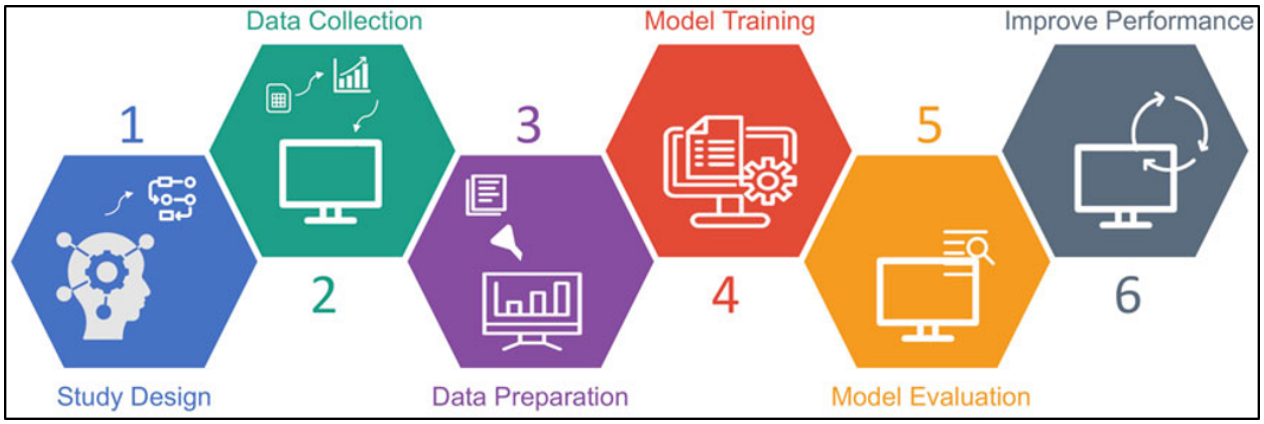

FIGURE 2.3 : Machine Learning Workflow.

This figure has been adapted from Peng et al. (2021), with the publisher's permission.

Having outlined the fundamentals of ML, we now turn to its application in plasma physics, where its capacity to model complex, nonlinear, and data-rich systems offers powerful tools for tackling the field's multi-scale and computational challenges.

2.2 Emergence of Machine Learning in plasma physics

The integration of ML into plasma physics has emerged as a promising approach for addressing the field's inherent complexity. Plasmas exhibit highly nonlinear, multiscale, and collective behaviors that make traditional modeling and simulation approaches both challenging and computationally demanding. ML offers data-driven methodologies that can complement physics-based models, providing new pathways for understanding, predicting, and controlling plasma dynamics. In computational plasma physics, ML is increasingly used to improve the accuracy, efficiency, and predictive power of simulations.

Its applications span a wide range of plasma regimes, from low-temperature plasmas in industry to high-temperature, laser-driven plasmas in fusion and high-energy-density physics. Table 2.1 compares these two types, highlighting the diverse physical conditions and modeling challenges they present. Despite their differences, both require tailored approaches for diagnostics, control, and simulation, needs that ML effectively addresses. As such, ML emerges as a unifying tool for tackling the complexity of both low- and high-temperature plasma environments.

Trieschmann et al. (2023) offers a comprehensive review of ML applications in low-temperature plasma modeling. Traditional plasma modeling faces challenges due to multi-scale, multi-physics dynamics and the complexity of plasma chemistry and plasma–surface interactions. The review describes ML methods, including deep learning, Gaussian Process Regression, physics-informed neural networks, and reduced-order modeling, that are increasingly applied to :

— accelerate or replace computationally expensive plasma simulations,

TABLE 2.1 : Low-temperature versus high-temperature plasmas.

Property	Low-Temperature Plasma	High-Temperature Plasma
Electron Temperature	1–10 eV (10,000–100,000 K)	100 eV – 10 keV (1–100 million K)
Ion Temperature	$\sim$0.1–1 eV (often much lower than electrons)	Comparable to electron temperature
Equilibrium State	Non-thermal equilibrium	Often thermal equilibrium
Ionization Level	Partially ionized	Fully or nearly fully ionized
Typical Density	10^{8}–10^{13} cm^{-3}	10^{18}–10^{22} cm^{-3}
Examples	Glow discharges, plasma jets, plasma medicine	Tokamaks, inertial confinement fusion, laser-produced plasmas
Applications	Surface treatment, etching, sterilization	Fusion energy, astrophysics, high energy density physics

— predict plasma chemistry, surface interactions, and sputtering yields,

— enable data-driven optimization and process control.

Since computational demands have long been a bottleneck in plasma research, Faraji et al. (2025) demonstrated how ML can significantly accelerate simulations while preserving accuracy. In high-energy-density physics, where plasmas experience extreme conditions, ML plays a crucial role in managing and interpreting large datasets. Likewise, Döpp et al. (2023) and Roussel et al. (2024) highlight ML-driven strategies in laser-plasma physics that facilitate experimental design, inverse problem-solving, and real-time optimization of plasma conditions. The studies emphasize the synergy between ML and high-performance computing to handle the complex and vast datasets generated in laser-plasma experiments. Furthermore, Gonoskov et al. (2019) employ ML, specifically artificial neural networks, to bridge theory and experiment in laser-plasma physics by analyzing high-dimensional spectral data. Their approach reduces the computational cost of parameter sweeps by using ML models as surrogate predictors, enabling faster theory validation and experimental optimization. These advancements collectively demonstrate that ML is not only enhancing computational plasma physics but also enabling new avenues for real-time control and predictive modeling of complex experimental plasma behaviors. In addition, Hatfield et al. (2021) discussed the data-driven future of high-energy-density physics, illustrating how ML models can rapidly identify complex interactions within large datasets. This capability is relevant for enhancing the predictive accuracy of simulations under extreme physical conditions. Hatfield et al. (2021) reveal that ML has already facilitated new insights into plasma turbulence, a critical factor affecting confinement in fusion devices.

The growing interest in ML applications within plasma physics is further evidenced by dedicated special issues and collaborative research topics. For example, the *Journal of Plasma Physics* published a special issue titled "Machine Learning for Plasma Physics and Fusion Energy" in 2022, collecting numerous contributions that explore the integration of ML, data science, and AI in plasma research. Similarly, *Frontiers in Physics* launched a research topic on "Machine Learning

in Plasma Physics, Chemistry, and Processing," aiming to disseminate the latest developments in ML methods across these disciplines. In addition to this, conferences such as the *Laser-Plasma Acceleration Workshop* or *Frontiers in Optics* series have included ML-dedicated sessions.

As ML continues to establish itself as a transformative tool in both low- and high-temperature plasma research, its relevance becomes particularly pronounced in the study of laser-plasma interactions. These interactions represent some of the most complex and nonlinear regimes in plasma physics, where traditional modeling struggles with the interplay of relativistic dynamics, steep gradients, and ultrafast processes. In this context, the following section delves into the intrinsic complexity of laser-plasma interactions and sets the stage for understanding how ML can help unravel their rich physical behaviors.

2.3 Complexity of laser–plasma interactions

High-intensity laser–plasma interactions involve a broad spectrum of nonlinear processes that dictate how energy is absorbed, transported, and converted into energetic particles. Understanding these dynamics is crucial in high-energy-density physics encompassing fields such as laser-plasma acceleration and inertial confinement fusion (ICF).

2.3.1 Plasma heating and hot electron generation

When an ultra-intense laser pulse strikes a target, plasma electrons absorb energy through several mechanisms that determine subsequent plasma behavior. These mechanisms generate hot electron populations that drive secondary plasma phenomena, including sheath formation and ion acceleration in TNSA, for example. A detailed discussion of the electromagnetic foundations of these processes is provided in Section 1.3.7.

2.3.2 Nonlinear propagation and self-focusing

As intense laser pulses propagate through plasma, nonlinear optical effects reshape the beam and influence energy deposition :

— **Relativistic self-focusing** arises because the oscillatory motion of electrons in the laser field increases their effective mass, reducing the local plasma frequency and increasing the refractive index in the beam core (Ritchie, 1994; Sprangle et al., 2007).

— **Ponderomotive self-focusing** occurs when the ponderomotive force expels electrons from high-intensity regions, forming a low-density channel that further guides the laser (Osman et al., 1999).

Together, these effects can increase on-axis intensity but also trigger filamentation and beam breakup, resulting in uneven energy deposition and reduced stability. Controlling self-focusing is therefore critical for stable particle acceleration. Approaches include using preformed plasma channels, multi-beam configurations, or applying external magnetic fields (Andreev et al., 1997; Hafizi et al., 2000).

Theoretical and computational models such as paraxial approximations (Hora, 1969), PIC simulations (Naseri et al., 2010), and source-dependent expansion methods (Choobini et al., 2024) have mapped out the regimes and thresholds of self-focusing. Figure 2.4 illustrates how the ponderomotive force shapes spot-size evolution under varying laser intensities.

2.3.3 Parametric and collective instabilities

Beyond self-focusing, intense laser–plasma interactions are often dominated by parametric instabilities that redistribute energy and limit coupling efficiency :

— **Stimulated Raman Scattering (SRS)** couples the laser to electron plasma waves, producing lower-frequency scattered light and hot electrons that can preheat the target, thereby degrading acceleration efficiency (Salcedo et al., 2003).

— **Stimulated Brillouin Scattering (SBS)** couples laser light to ion acoustic waves, causing backscattering and depletion of the driving pulse, which lowers the energy available for plasma heating and particle acceleration (Giulietti et al., 1999).

— **Two-plasmon decay (TPD)** occurs when a photon decays into two plasmons, generating hot electrons with properties often suboptimal for acceleration, and contributing to premature target heating (Yao et al., 2024a; Ruskov et al., 2024; Lian et al., 2025).

Collectively, these instabilities can remove a substantial fraction of the incident laser energy, fragment the beam, and induce uncontrolled plasma heating. The resulting beam breakup and phase distortions reduce the peak intensity at focus, while excessive or nonuniform electron heating alters plasma profiles, preheats the target, and weakens the density gradients needed for efficient acceleration.

Their impact is particularly severe in the relativistic and near-critical density regimes, where SRS and SBS growth is enhanced (Weber et al., 2008; Moreau et al., 2017). In these regimes, the two processes can be coupled through shared plasma wave dynamics and density modulations, meaning that suppressing one instability can unintentionally amplify another. Next-generation multi-PW laser systems will operate closer to these instability thresholds due to higher intensities, longer interaction lengths, and higher repetition rates. As a result, effective mitigation strategies such as plasma tailoring, multi-beam geometries, or the application of strong external magnetic fields must account for their interplay. A comprehensive understanding of these coupled instabilities is therefore essential for ensuring stable, efficient energy coupling in future high-energy laser facilities.

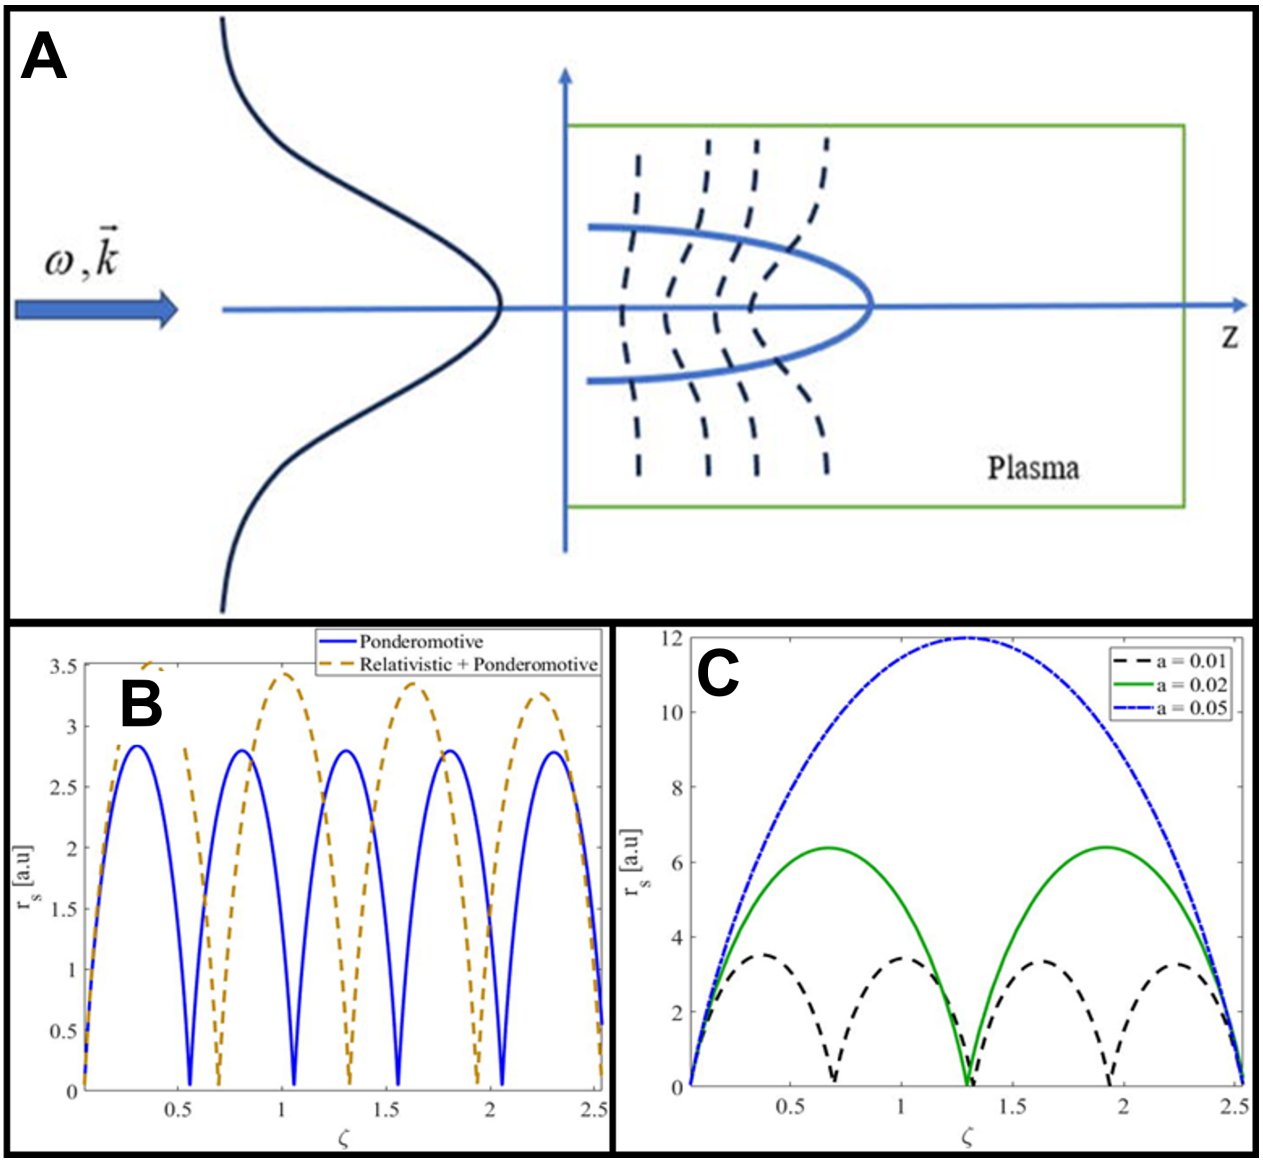

FIGURE 2.4 : Self-focusing of a laser beam in magnetized plasma.

(A) Schematic of self-focusing of a laser beam in magnetized plasma. z is the propagation axis; the full lines represent the norm of the wave vector $|\mathrm{k}|$ before (black) and in the plasma (blue); the dashed lines represent the wavefront. (B) Impact of the ponderomotive force on the amplitude of spot size of the laser beam vs the normalized distance of propagation $(\zeta = z/z_R)$ for $\lambda = 800\,\mathrm{nm}$, $\gamma = 1.1$, $\omega_p = 2\,\mathrm{THz}$, $B_0 = 100\,\mathrm{T}$, $\sigma = +1$, and $a = 0.01$. (C) Variations in the spot size of the laser beam vs the normalized distance of propagation (ζ) for different laser intensities and for $\lambda = 800\,\mathrm{nm}$, $\gamma = 1.1$, $\omega_p = 2\,\mathrm{THz}$, $B_0 = 100\,\mathrm{T}$, and $\sigma = +1$. These figures have been adapted from Choobini et al. (2024), with the publisher's permission.

2.3.4 Experimental diagnostics and modeling

Significant progress in diagnostic techniques and simulations has deepened our understanding of these interactions :

— **PIC simulations** provide kinetic-level insights into phenomena such as filamentation, hot-electron generation, and sheath evolution, and can incorporate coupled atomic and plasma processes (Copplestone et al., 2019).

— **Thomson scattering**, which involves the scattering of laser light by plasma electrons, allows non-invasive measurements of electron temperature and density in dense plasmas (Kaloyan et al., 2021, 2022).

— **Proton radiography** has visualized evolving electromagnetic fields and plasma bubble dynamics in ultra-intense interactions (Mackinnon et al., 2004; Romagnani et al., 2005).

These tools have clarified how self-focusing, filamentation, and parametric instabilities arise and interact, guiding stabilization strategies and informing predictive models.

2.4 High-dimensional parameter space in laser–plasma acceleration

Laser-driven ion acceleration (LDIA) emerges from the coupled influence of laser parameters, target properties, pre-plasma conditions, and environmental factors. Optimizing one parameter inevitably affects others, creating a high-dimensional landscape that governs the laser-driven beam quality.

2.4.1 Laser parameters

Laser beam energy, pulse duration, temporal contrast, and focusing geometry (wavefront) are strongly interdependent and jointly determine the efficiency, stability, and spectral quality of LDIA

— **Pulse energy** sets the total energy available for generating hot electrons and accelerating ions in TNSA and RPA (Esirkepov et al., 2004; Fuchs et al., 2006). Increasing energy generally enhances ion yield and maximum energy but without adequate temporal contrast, excessive energy can pre-expand targets and reduce coupling efficiency (McKenna et al., 2006; Ceccotti et al., 2007; Kim et al., 2024). Efficient use of high-energy pulses therefore requires contrast enhancement techniques such as plasma mirrors or cross-polarized wave generation.

— **Pulse duration** defines the effective acceleration timescale. Sub-picosecond pulses produce higher peak intensities and narrower energy spectra, suppressing undesired interactions and minimizing target expansion (Oishi et al., 2005; Fuchs et al., 2006; Fourmaux et al., 2013; Simpson et al., 2021b). However, pulses that are too short deposit insufficient energy to sustain a stable sheath (Mackinnon et al., 2002), while overly long pulses promote premature plasma blow-off. Optimal durations thus balance peak intensity with sustained electron heating.

— **Focal geometry and wavefront quality** govern the transverse scale of energy deposition, rear-sheath uniformity, and ultimately the divergence and cutoff energy of TNSA proton beams. Afshari et al. (2020) showed that replacing a tightly focused Gaussian spot ($5\ \mu$m FWHM) with a smoothed top-hat profile ($35\ \mu$m) on the *PHELIX* laser reduced proton divergence from $\sim 16^\circ$ to $\sim 2^\circ$ and boosted total yield, due to a flatter, wider sheath field. Wang et al. (2025) demonstrated that Laguerre–Gaussian beams both increase proton energy and halve the angular spread compared to Gaussian beams. Real-time adaptive optics further enhances wavefront control : for example, the *ARTAO* system introduced by Ohland et al. (2025) stabilizes wavefronts on a shot-to-shot basis under vacuum, using fast wavefront sensing, a closed-loop controller, and a bimorph deformable mirror. This approach is especially advantageous for high-repetition-rate systems, where thermal drift and laser fluctuations would otherwise compromise reproducibility. Adaptive optics thus not only mitigate aberrations but also actively enable fine-tuning of focal conditions, opening new degrees of freedom for controlling laser–plasma coupling and improving laser-plasma acceleration performance (Fourmaux et al., 2008; Lin et al., 2019; Oumbarek Espinos et al., 2023; Tang et al., 2024).

While laser parameters set the fundamental conditions for generating and shaping the accelerating fields, their influence is inseparable from the state of the target at the time of interaction. In particular, the temporal contrast, focal geometry, and wavefront quality discussed above not only determine the intensity profile at focus but also govern the formation of pre-plasma and the degree to which the target surface is modified before the arrival of the main pulse. Consequently, to fully optimize laser–plasma coupling in LDIA, these beam characteristics must be considered alongside target properties and pre-plasma conditions, which together dictate the efficiency, stability, and spectral quality of the accelerated particle beams.

2.4.2 Target properties and pre-plasma conditions

The impact of target material composition, thickness, and density on the efficiency of TNSA and RPA has already been thoroughly discussed in Section 1.3.14. In LWFA using gas jets, the plasma density profile — shaped by the gas type, backing pressure, and nozzle geometry — plays a crucial role in determining the wakefield strength, electron injection conditions, and beam quality, making precise control of gas jet parameters essential for efficient acceleration. Several studies have dealt with this reality : Lorenz et al. (2019); Zhou et al. (2021); Cobo et al. (2024).

In LDIA, especially in the TNSA regime, pre-plasma conditions are crucial in shaping both the efficiency and quality of the accelerated ion beams. Among the key factors influencing these conditions are pre-pulse formation and the pre-plasma scale length, both of which significantly impact laser absorption and energy transfer dynamics.

A pre-pulse, defined as a lower-intensity laser pulse arriving before the main high-intensity pulse, can ionize the target surface and form a pre-plasma that substantially modifies the subsequent laser–target interaction. The presence of pre-plasma impacts laser absorption, hot-electron generation, and the sheath electric fields that drive ion acceleration.

Controlled pre-plasma formation has been shown to significantly enhance proton acceleration. Micrometer-scale pre-plasmas can produce substantial increases in proton cut-off energy and beam charge; in some cases, three-fold energy gains have been observed due to improved coupling of laser energy into hot electrons (Gizzi et al., 2021). However, excessive pre-plasma formation can be detrimental, leading to increased beam divergence and reduced proton energy, underscoring the need for precise control of pre-pulse parameters. Although several experimental and simulation studies support these trends (Hadjisolomou et al., 2020; Gizzi et al., 2021; Kim et al., 2022; Strehlow et al., 2022; Keppler et al., 2022), the optimal pre-plasma conditions remain strongly dependent on laser parameters, target properties, and the specific acceleration regime.

A key parameter is the pre-plasma scale length, the distance over which the plasma density rises from underdense to critical, as it governs electron heating, laser absorption, and ultimately the efficiency of ion acceleration. Vladisavlevici et al. (2025) show that increasing the scale length can significantly enhance absorption into hot electrons, sometimes up to 80–90%, because the laser interacts more effectively with near-critical regions. This enhanced coupling can raise the maximum proton cutoff energy, but only up to an optimum scale length : for example, about 1 μm in thin titanium foils and $\sim$3 μm in thick CH targets, beyond which electron dilution and weaker sheath formation reduce acceleration efficiency. Thus, while longer scale lengths can improve laser–electron coupling, efficient TNSA requires balancing this gain against the degradation of the rear-side sheath field, leading to a target- and geometry-dependent optimum rather than a simple monotonic trend.

To mitigate the detrimental effects of pre-pulses, techniques such as plasma mirrors, cross-wave polarizers, and saturable absorbers have been developed. Plasma mirrors enhance contrast by reflecting only the high-intensity main pulse, allowing low-intensity pre-pulses to pass without interaction. Cross-wave polarizers use polarization filtering to suppress pre-pulses while transmitting the main pulse. Saturable absorbers block low-intensity light but become transparent at high intensities, effectively filtering out pre-pulses. This enhancement ensures that the main pulse interacts with a well-defined target surface, thereby optimizing the conditions for efficient ion acceleration.

2.4.3 Environmental factors

Residual gas pressure and applied magnetic fields can also shape LDIA :

— **Low ambient pressure** ($< 10^{-2}$ mbar) minimizes unwanted interactions and improves conversion efficiency (Snyder et al., 2020).

— **External magnetic fields** (up to kT-level) confine hot electrons, enhance sheath fields, and collimate ion beams (Sarri et al., 2012; Ferri et al., 2019a; Weichman et al., 2020; Khan et al., 2024).

Careful environmental control can thus complement laser and target optimization.

2.4.4 Conclusion

Laser–plasma interactions are inherently multi-parametric and nonlinear, involving coupled processes of electron heating, self-focusing, and collective instabilities. Efficient LDIA requires a joint optimization of laser parameters, target properties, pre-plasma, and environmental conditions.

Modern diagnostics, high-fidelity simulations, and ML-based optimization are increasingly essential to explore this high-dimensional space, mitigate deleterious instabilities, and guide next-generation high-repetition-rate laser–plasma accelerators toward reliable, high-performance operation.

2.5 Machine Learning techniques for laser–plasma physics optimization

Recent advancements in ML have further augmented real-time optimization capabilities in TNSA experiments. The works of Döpp et al. (2023) and Roussel et al. (2024) provide a comprehensive overview of how ML techniques are transforming laser-plasma physics, particularly in areas like laser-plasma acceleration and ICF. With advancements in high-power laser technology, researchers can now collect extensive datasets from both experiments and simulations. The authors discuss the application of various ML methods, including Supervised Learning for modeling and prediction, Unsupervised Learning for data exploration, and optimization algorithms for enhancing experimental setups. They highlight how these data-driven approaches enable more efficient analysis, improved predictive capabilities, and the potential to uncover new physical insights in laser-plasma interactions.

2.5.1 Artificial neural networks

Artificial neural networks (see Figure 2.5) are a class of ML models inspired by the structure and functioning of the human brain. They are the backbone of many technologies we use today, from facial recognition and self-driving cars to language translation and *ChatGPT* itself. At their core, neural networks are function approximators. They learn to map inputs to outputs by adjusting internal parameters through a process called training. The concept of a neural network originates from neuroscience. The human brain consists of billions of neurons connected by synapses. Each neuron receives signals from others, processes them, and passes the signal along if certain conditions

are met. Artificial neural networks take inspiration from this architecture. While vastly simplified, they mimic the brain's way of learning patterns from data.

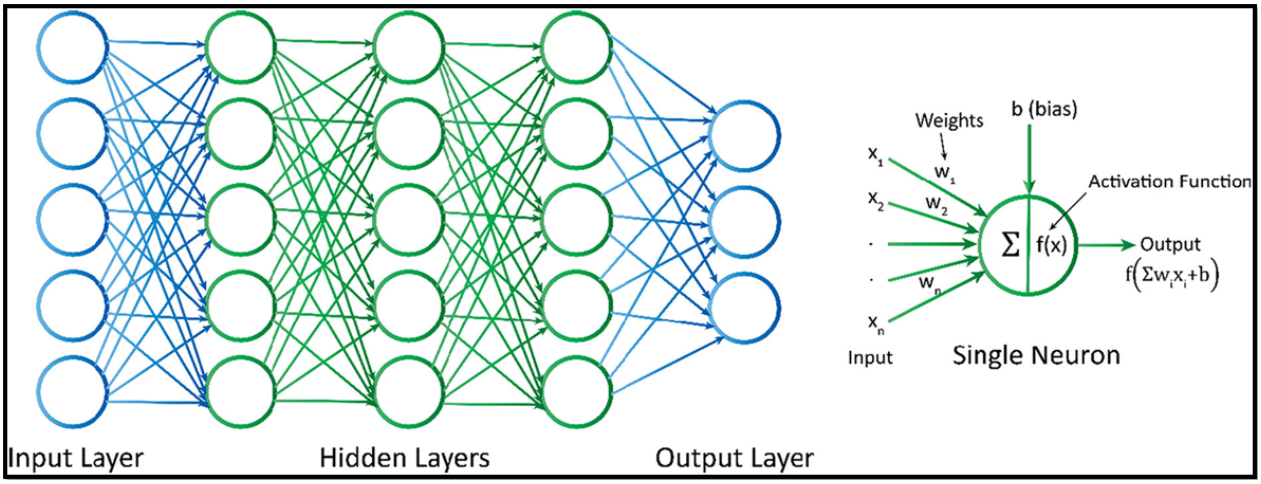

FIGURE 2.5 : Feedforward neural network and single neuron model.

This figure has been adapted from Baduge et al. (2022), with the publisher's permission.

A basic neural network consists of three types of layers :

- **Input layer :** Receives the raw data (e.g., pixels of an image, words in a sentence).
- **Hidden layers :** These perform transformations of the input. The more hidden layers there are, the deeper the network, hence the term deep learning.
- **Output layer :** Produces the final result (e.g., classification label or predicted value).

Each layer contains nodes (or neurons), and each node performs a simple mathematical operation : it computes a weighted sum of its inputs, adds a bias, and passes the result through a nonlinear activation function. This allows the network to model complex, nonlinear relationships.

Neural networks learn by adjusting their weights and biases to minimize a loss function, which measures the error between the network's predictions and the actual outputs. This optimization (minimizing a loss function) is done using a process called backpropagation, combined with gradient descent. Here's how it works :

1. Data is fed into the network to make predictions.
2. The loss is calculated.
3. The gradient of the loss is computed with respect to every weight in the network.
4. Each weight is updated slightly in the direction that reduces the error.

This process is repeated over many cycles, called epochs, gradually improving the network's performance.

Different tasks require different network architectures :

— **Feedforward Neural Networks (FNN) :** The simplest type; data flows in one direction (Figure 2.5).
— **Convolutional Neural Networks (CNN) :** Used for image processing; they recognize spatial hierarchies in data.
— **Recurrent Neural Networks (RNN) :** Designed for sequential data like time series or language; they have memory of previous inputs.
— **Transformers :** A newer architecture that has revolutionized language processing. Unlike RNNs, transformers use attention mechanisms and process data in parallel.

Neural networks are powerful because they can automatically extract features from raw data and approximate complex functions with high accuracy, often beyond what traditional models can do. They are universal function approximators, meaning that with enough data and the right architecture, they can model virtually any input-output relationship. This ability has led to breakthroughs in :

— Medical diagnosis from imaging,
— Natural language processing,
— Physics-informed modeling,
— Autonomous vehicles.

Despite their success, neural networks come with limitations :

— **Data hungry :** They often require large datasets to generalize well.
— **Opaque :** They are typically black-box models, hard to interpret.
— **Computationally expensive :** Training deep networks demands significant hardware and energy.
— **Overfitting :** Without proper regularization, they can memorize instead of generalize.

Thus, there is an ongoing push toward making neural networks more efficient, interpretable, and robust.

Building on these foundational concepts, neural networks have been increasingly applied to model and optimize laser–plasma interactions. Recent studies demonstrate how their capacity to learn complex nonlinear mappings between experimental or simulated inputs and beam characteristics can significantly accelerate parameter exploration and reduce reliance on computationally expensive PIC simulations or extensive experimental campaigns. Schmitz et al. (2023) modeled a liquid leaf target TNSA experiment using 2D PIC simulations combined with deep neural networks to predict proton spectra from laser and target parameters. The trained surrogate model reproduced PIC results with high accuracy, enabling rapid parameter scans and reducing the need for extensive simulations or experimental shots. This data-driven approach offers a fast, reliable tool for optimizing laser–plasma experiments, particularly in high-repetition-rate facilities. McQueen et al. (2025) recently introduced a synthetic diagnostic framework based on deep neural networks, capable of predicting laser-driven proton spectra from experimental input parameters. This approach facilitates high repetition-rate optimization, achieving accurate proton beam characteristics prediction

within approximately 700 laser shots. Djordjević et al. (2021b) investigate the application of deep learning techniques to optimize ion acceleration driven by short-pulse lasers. The authors employ neural networks as surrogate models to efficiently map extensive parameter spaces, enabling the identification of optimal conditions for ion acceleration. A significant finding from this study is the observed dependence of ion energy on the pre-plasma gradient length scale, highlighting the intricate dynamics of the acceleration process. The methodologies presented in this research offer a promising approach to reducing computational costs associated with traditional simulation techniques. The same group (Djordjević et al., 2021a) developed a stacked convolutional and recurrent neural network (CNN-RNN) as a surrogate model to predict ion acceleration time in short-pulse, LDIA. The model is trained on a large ensemble of PIC simulations, incorporating time dependencies for improved regression. Pretraining on lower-fidelity 1D analytical models enhances the learning of high-fidelity simulation data. The study explores how laser and plasma parameters affect acceleration time, providing a data-driven approach to optimizing ion acceleration dynamics while reducing computational costs. The authors of Mariscal et al. (2024) developed a neural network with uncertainty quantification to predict the number of hot electrons generated during laser–target interactions from various pulse-shaping parameters. To validate their approach, they compared two models : one trained exclusively on experimental data and another on ensemble PIC simulations. Both predicted a potential increase in hot electron generation of approximately 12–18%; however, experiments did not show a statistically significant enhancement. This discrepancy mainly stems from the inherent gap between idealized simulations and real experimental conditions. Simulations assume perfectly controlled parameters, identical targets, precise laser energy and no hardware or diagnostic noise whereas experiments face inevitable fluctuations such as energy jitter, pre-pulses, beam pointing drifts and alignment errors. Moreover, simulations often explore a broader parameter space than what is experimentally accessible, capturing stronger theoretical trends that real systems cannot safely reproduce. Finally, while experimental data are affected by $\sim$20–30% shot-to-shot noise and detector artifacts, simulations exhibit $\leq$1% numerical variation, yielding smoother and more correlated results. Together, these factors explain the differing trends observed between simulation- and experiment-trained models.

2.5.2 Support Vector Regression

Support Vector Regression (SVR) (Smola et al., 2004; Basak et al., 2007; Awad et al., 2015) is an ML technique adept at modeling complex nonlinear relationships. SVR operates by mapping input data into a higher-dimensional feature space using kernel functions, enabling it to capture intricate patterns within the data (Zhang, 2023). The goal is not to assign a class but to predict a real number, such as temperature or energy output. This means that SVR allows small errors without penalty, only focuses on data points outside this tolerance range (called support vectors), and tries to make the model as simple as possible while remaining within acceptable error.

To visualize the SVR principle, imagine drawing a tube around a predicted curve. This tube has a certain width, defined by ϵ. SVR tries to find the curve such that :

— Most of the training points fall inside the tube.

— The model is as simple and smooth as possible (*i.e.*, it avoids overfitting).

— Only points outside the tube (those with error greater than ϵ) influence the model's parameters.

This leads to a sparse solution : the final model only depends on a subset of the data. Those points lying outside the tube are known as support vectors. This principle is illustrated in Figure 2.6 in 2D where $\epsilon =$ Maximum margin.

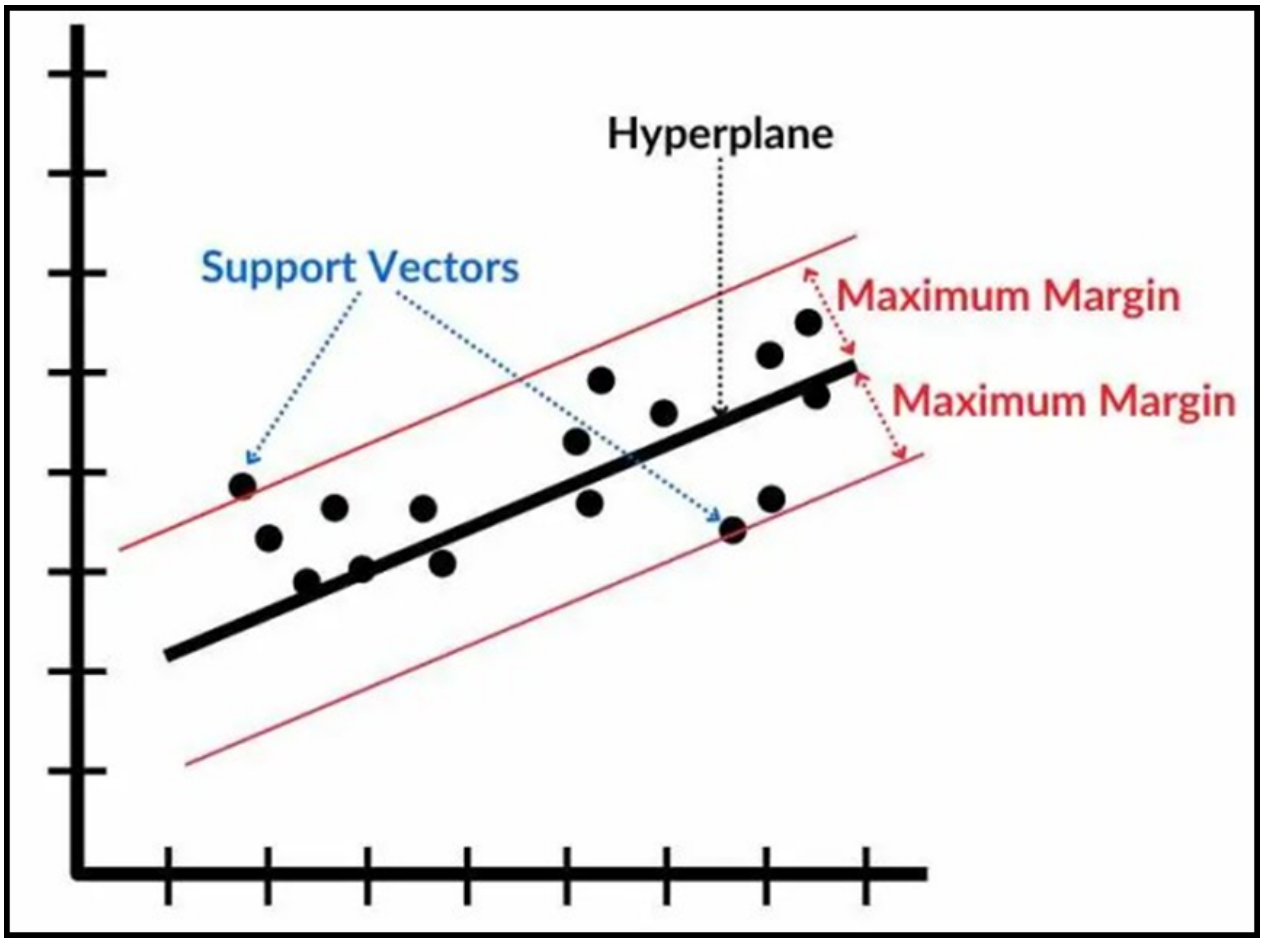

FIGURE 2.6 : Support Vector Regression representation.

This figure has been adapted from Spot Intelligence (2024), with the publisher's permission.

One of the strengths of SVR is that it can be extended to handle nonlinear relationships. Instead of fitting a straight line or flat plane, SVR can project the data into a higher-dimensional space where a linear model becomes a nonlinear model in the original space.

Common kernels include :

— **Linear kernel :** for simple problems.

— **Polynomial kernel :** for capturing curved trends.

— **Radial Basis Function (RBF/Gaussian) kernel :** for flexible, smooth fits.

Compared to linear regression, SVR is more robust to outliers thanks to its margin-based formulation. Relative to neural networks, SVR typically requires fewer data points and is less prone to overfitting, although it can become computationally expensive for very high-dimensional or large-scale datasets. Here, high dimensionality refers to cases where the number of input features exceeds several tens of correlated variables, while large datasets correspond to training sets comprising more than a few thousand experimental shots or simulation samples. In a recent study, Desai et al. (2024) compared three ML models — SVR, Gaussian Process Regression (GPR), and a three-hidden-layer neural network — for optimizing and predicting key outputs in laser-driven proton acceleration. Using a synthetic dataset derived from a modified Fuchs et al. (2006) model that incorporates experimental realism such as focal position variation and noise, the authors optimized parameters including laser intensity, target thickness, and target position relative to focus. The performance metrics evaluated were maximum proton energy, total kinetic energy, average energy, and laser-to-proton energy conversion efficiency, with training sets of up to 20,000 synthetic data points. Results showed that SVR provided a favorable balance between accuracy and computational efficiency, GPR offered strong predictive performance with robust uncertainty quantification, and neural networks required larger datasets to achieve comparable accuracy. A key limitation of the SVR results reported by Desai et al. (2024) is that the model was trained exclusively on synthetic data, which is smoother and less feature-rich than real experimental datasets. Therefore, SVR does not provide inherent uncertainty quantification, which is valuable in noisy, shot-limited experiments. Furthermore, in optimization tests, SVR exhibited systematic biases, favoring thinner targets and focal offsets that deviated from the ground-truth model. These factors, combined with the lack of validation on real experimental data, limit the immediate transferability of the reported SVR performance to complex, high-repetition-rate laser–plasma experiments.

The efficacy of SVR in modeling laser-plasma interactions can be attributed to several factors :

— **Nonlinear modeling capability :** SVR's use of kernel functions allows it to capture complex, nonlinear relationships inherent in laser-plasma dynamics.

— **Robustness to overfitting :** By focusing on a subset of critical data points (support vectors), SVR maintains generalization performance, even with limited data.

— **Computational efficiency :** SVR's optimization framework ensures efficient training, making it suitable for real-time applications where rapid predictions are essential.

Despite its advantages, SVR presents several limitations when applied to LDIA :

1. **Scalability issues :**

   — SVR's computational cost increases significantly with large datasets, as its training complexity is $\mathcal{O}(n^2)$ to $\mathcal{O}(n^3)$, where n is the number of data points.

— This makes SVR less practical for handling high-dimensional datasets generated from large-scale PIC simulations or experimental data with thousands of input features.

2. **Limited extrapolation ability :**
   — Unlike deep learning models, which can generalize better to unseen experimental conditions, SVR struggles with extrapolating beyond the parameter range of its training data.
   — This limits its usefulness in highly dynamic, exploratory experimental setups where unexpected plasma behavior may occur.
3. **Kernel selection sensitivity :**
   — The choice of kernel function (e.g., polynomial, Gaussian, RBF) heavily influences the accuracy of SVR predictions.
   — Selecting an inappropriate kernel may lead to poor model performance, requiring extensive hyperparameter tuning to optimize.
4. **Difficulty in handling noisy experimental data :**
   — Laser-plasma experiments inherently produce noisy and sparse datasets due to fluctuations in laser intensity, target surface conditions, and pre-plasma formation...
   — SVR, particularly with complex kernels, may struggle to filter out experimental noise effectively compared to deep learning methods that can learn noise-resilient representations.

Several strategies can mitigate these limitations :

— **Cross-validation and regularization :** tuning the penalty parameter C and kernel width γ helps control model complexity and prevent overfitting to noisy data.
— **Data augmentation and feature scaling :** expanding and normalizing the training space improves kernel performance in high-dimensional parameter domains.
— **Hybridization with probabilistic models :** combining SVR with Gaussian Process Regression (SVR–GPR) can provide uncertainty-aware predictions.

Hybrid approaches also offer promising paths forward. For instance, SVR–Neural Network ensembles leverage SVR's robustness for coarse trend prediction and use neural networks to refine residual errors or extract nonlinear correlations. Alternatively, SVR–Random Forest frameworks can exploit the interpretability of tree-based models to guide feature importance and adaptive kernel weighting. Such hybridization can enhance both accuracy and generalization, making SVR-based methods more suitable for noisy, high-dimensional laser–plasma experiments.

While SVR remains a powerful and interpretable tool for modeling nonlinear dependencies in laser–plasma systems, its performance can degrade when confronted with highly noisy, nonstationary, or strongly multidimensional experimental data. These limitations are not prohibitive but highlight the importance of careful kernel selection, regularization, and dataset design to ensure physical consistency and generalization. In practice, SVR performs best when used alongside

complementary methods, such as Gaussian Process Regression for uncertainty quantification or neural networks for feature extraction, within hybrid or ensemble frameworks. Such approaches can combine SVR's robustness and interpretability with the adaptability of deep learning, offering a balanced path toward accurate and reliable modeling of complex laser–plasma dynamics.

2.5.3 Random Forests

Among the various supervised learning methods used to model complex relationships between input parameters and physical observables, Random Forest (RF) models have emerged as one of the most robust and practical approaches. Their good generalization capabilities, low sensitivity to noise, and ability to handle multivariate, non-linear dependencies make them particularly suited for laser-driven particle acceleration studies, in which experimental data are often limited, noisy, and affected by shot-to-shot fluctuations.

RF is an ensemble-learning algorithm proposed by Breiman (2001), based on aggregating multiple decision trees. Each tree is trained on a randomly resampled dataset (bootstrap sampling), and at each node a random subset of features is used for splitting. For regression tasks, the final prediction is the average output $\hat{y}$ of all N_{trees} trees :

Equation 2.1 : Random Forest regression prediction

$$\hat{y}(\mathbf{x}) = \frac{1}{N_{\text{trees}}} \sum_{i=1}^{N_{\text{trees}}} T_i(\mathbf{x}) \tag{2.1}$$

where T_i is the i-th decision tree and $\mathbf{x}$ is the feature vector (e.g., laser parameters...). Two sources of randomness - bootstrap sampling and random feature selection — reduce correlation among trees and thereby lower the variance of the aggregated estimator (Hastie et al., 2008).

Each tree is grown by recursively partitioning the feature space to minimize an impurity measure (e.g., mean squared error for regression). Hyperparameters include the number of trees (N_{trees}), maximum depth, minimum samples per split/leaf, and the number of features sampled per split. In practice, RFs are remarkably stable to these choices; performance typically improves monotonically with N_{trees} until saturation (Breiman, 2001).

Laser-driven ion acceleration (e.g., TNSA) depends on many coupled parameters : wavefront structure, pulse duration and spectral phase, target thickness/material, pre-plasma evolution, focusing geometry, and contrast. These dependencies are (i) highly nonlinear, (ii) non-invertible analytically, (iii) sensitive to initial conditions, and (iv) often observed through sparse, noisy datasets.

A particularly useful property is feature importance analysis, which quantifies how much each input contributes to prediction quality. In practice, this helps identify a subset of the most influential control knobs (e.g., a handful of DM actuators or specific spectral-phase terms), reducing search

TABLE 2.2 : Experimental challenges vs. Random Forest strengths.

Experimental Challenge	RF Strength
Noisy, sparse data (few 10s–100s of shots)	Bagging/averaging yields noise-robust predictions (Breiman, 2001)
Nonlinear, high-order feature interactions	Trees capture complex, nonparametric relations
Need for interpretability (what knobs matter?)	Feature importance (MDI/Permutation) (Liaw et al., 2002)
Limited need for heavy hyperparameter tuning	Good defaults; graceful scaling with N_{trees}

dimensionality and informing subsequent optimization (Liaw et al., 2002). Table 2.2 shows experimental challenges and strengths of the RF regression.

Illustrative example

To illustrate (Figure 2.7) the RF process in its simplest form, consider a toy example in which we wish to model the dependence of the proton maximum energy $\mathcal{E}_{p,\max}$ on only two experimental inputs : (i) x_1 the deformable mirror defocus actuator voltage and (ii) x_2 the spectral phase term β_2 (group-delay dispersion). Suppose we have 10 experimental shots :

TABLE 2.3 : Illustrative toy dataset.

Shot #	x_1 [V]	x_2 [fs^2]	$\mathcal{E}_{p,\max}$ [MeV]
1	-40	-200	1.8
2	-20	-150	2.5
3	0	-100	2.9
4	20	-50	3.4
5	40	0	3.1
6	60	50	2.7
7	80	100	2.4
8	100	150	2.1
9	120	200	1.9
10	140	250	1.6

Training. The RF trains many decision trees on bootstrap-resampled versions of Table 2.3. Each tree sees a slightly different dataset and selects either x_1 or x_2 at each split, forming a set of simple rules.

Prediction. For a new configuration $(x_1, x_2) = (30\ \text{V}, -70\ \text{fs}^2)$, each tree returns a slightly different estimate $T_i(x_1, x_2)$; the forest prediction is the average : $\hat{\mathcal{E}}_{p,\max} = \frac{1}{N_{\text{trees}}} \sum_{i=1}^{N_{\text{trees}}} T_i(30, -70) \approx 3.25\ \text{MeV}$.

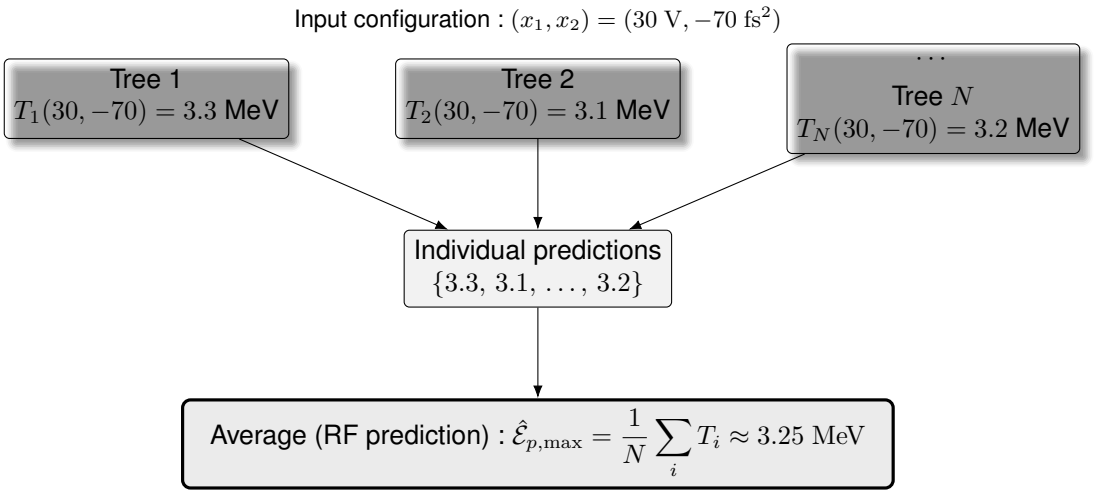

FIGURE 2.7 : Random Forest prediction using multiple decision trees.

Illustration of the Random Forest inference process. Each decision tree produces an independent prediction, and the final output is obtained by averaging across all trees to reduce variance and improve robustness.

Feature importance. The RF also indicates which feature matters most. In this toy case, a typical outcome is :

Feature	Importance
x_2 (GDD)	0.70
x_1 (Defocus)	0.30

suggesting that dispersion influences $\mathcal{E}_{p,\max}$ more than defocus in this simple two-parameter system.

2.5.4 Limitations and Complementarity

Despite their strengths, RFs have limitations :

— **Extrapolation :** Predictions are trustworthy within the training domain; RFs do not extrapolate physically beyond observed data.

— **Non-smoothness :** The predictor is piecewise constant across leaf regions; when a smooth functional surrogate is needed, Gaussian processes are preferable (Snoek et al., 2012).

— **Model size :** Large forests can be memory-heavy.

Practical Tips

— **Start simple :** Use $N_{\mathrm{trees}} = 300$–500, moderate max depth (e.g., 8–14), and enable out-of-bag (OOB) error to track generalization without extra validation splits.

— **Stability checks :** Report both impurity-based and permutation-based importances; inspect collinearity and cluster inputs that move together.

— **Physics awareness :** Constrain the search to physically safe/meaningful regions (e.g., avoid spectral holes; respect hardware voltage/current limits; account for contrast).

In the context of relativistic laser–plasma experiments, RFs have been successfully used as supervised regression tools to link experimental controls with plasma-beam observables. For example, in laser-wakefield acceleration, RF models were trained to predict the accelerated electron charge from laser-wavefront parameters and were shown to offer competitive accuracy and superior interpretability, revealing which Zernike modes most strongly impact beam generation (Lin et al., 2021). At larger scales, ensemble-based supervised learning is also cited as an effective approach in inertial ICF research, complementing deep-learning frameworks for surrogate modeling, prediction and diagnostic inference, particularly where datasets are limited and model interpretability remains valuable (Humbird et al., 2019). Taken together, these studies demonstrate that RFs are especially attractive for plasma physics because they handle high-dimensional nonlinear parameter spaces, tolerate experimental noise, and provide direct physical insight through feature-ranking capabilities that are difficult to obtain from purely optimization-oriented or black-box neural models.

Indeed, RFs offer a powerful and interpretable tool for laser-driven particle acceleration. Their robustness against noise, ability to capture nonlinear physics, and capability to reduce dimensionality make them well-suited for data-driven control of high-intensity laser–plasma sources.

2.5.5 Gaussian Process regression

While less well-known outside of research, Gaussian Process regression (GPR) is a gold standard for regression problems when interpretability, uncertainty estimation, and small-data performance are essential (Rasmussen, 2003; Quinonero-Candela et al., 2005). At its heart, GPR is a non-parametric, probabilistic model. Most models try to find one fixed formula that fits your data. GPR works differently; instead of guessing just one curve, it considers many possible curves that could match your data, each with its own probability. We can think of it as saying, “Here are all the shapes the data might take, and here’s how sure we are about each part of the shape.” In GPR, the predictions at different points are linked, so if you learn something in one region, it also changes what you believe in nearby regions. This means it doesn’t just give you numbers with error bars, it tells you how confident it is about the entire curve. Formally, a GP is defined by a mean function $m(x)$ and a covariance function (or kernel) $k(x, x')$:

Equation 2.2 : Definition of a Gaussian Process

$$f(x) \sim \mathcal{GP}\big(m(x),\, k(x, x')\big) \tag{2.2}$$

where x and x' are input points. Here, $\sim$ indicates that the random variable follows a certain distribution, in this case, a GP. This notation means that for any finite set of input points, the corresponding function values follow a multivariate normal distribution whose mean and covariance are given by evaluating $m(x)$ and $k(x, x')$ at those inputs. The kernel function $k(x, x')$ measures the similarity between inputs : it is largest when $x = x'$ and typically decreases as the distance between x and x' increases.

Choosing the GP Mean and Kernel

The mean encodes the large-scale trend you expect before seeing data :

— **By default :** $m(x) = 0$ (works well if data are centered or detrended).

— **We use non-zero mean** when a trend is known.

— **Tip :** If you subtract the average trend from your data (detrending), you can still safely use a zero mean.

The choice of kernel dictates how the model interprets correlations in the data. In practice, the kernel controls properties such as smoothness, correlation length, periodicity, and the presence of multiple characteristic scales. Several standard kernels are widely used :

— **Squared Exponential (RBF)** : $k(x, x') = \sigma_f^2 \exp\left(-\frac{(x-x')^2}{2\ell^2}\right)$, which produces very smooth functions and is often a reliable default;

— **Matérn kernel**, which introduces a tunable roughness parameter ν, allowing the model to represent less regular and more realistic experimental behaviors;

— **Periodic kernel**, well-suited for signals exhibiting repeating or oscillatory structure;

— **Linear kernel**, which captures global trends or drift in the data.

Because real systems rarely behave according to a single structure, kernels can be combined through addition or multiplication to model more elaborate behavior (e.g. additive long- and short-scale effects, or periodic fluctuations around a trend). The corresponding hyperparameters — including the length scale ℓ, the signal variance σ_f^2, and the noise variance σ_n^2 — are typically inferred by maximizing the marginal likelihood, ensuring that the chosen kernel makes the observed data statistically most plausible.

In most applications, an effective strategy is :

1. Begin with a simple prior, using a zero mean $m(x) = 0$ and either an RBF or Matérn kernel;
2. Optimize the kernel hyperparameters $(\ell, \sigma_f^2, \sigma_n^2)$ via marginal likelihood;
3. Evaluate residuals and learning curves to identify remaining structure; if patterns persist, refine the kernel by :
   — adding or multiplying a periodic kernel for oscillatory behavior,
   — summing multiple RBF kernels for multi-scale dynamics,

— introducing a non-zero mean when a clear baseline trend is known.

4. Validate the resulting model using held-out data or cross-validation to ensure generalization.

As a guideline, the mean should describe only the global trend, while the kernel accounts for local variations. In addition, keeping the mean function simple generally leads to more robust inference, since the kernel naturally carries most of the structural information. Finally, when working with experimental data, Matérn kernels often provide a better fit than the overly smooth RBF kernel.

One of the key strengths of GPR lies in its Bayesian interpretation :

— Rather than fitting a single curve to the data, GPR considers a range of possible functions consistent with the observed data and the prior defined by $m(x)$ and $k(x, x')$.

— It provides both predictions and uncertainty estimates at each input point, making GPR highly interpretable and particularly useful in scientific and exploratory modeling.

Illustrative Example : Sampling from a Gaussian Process Prior

To visualize how a GP defines a distribution over functions, we construct an illustrative example where functions are sampled from a GP prior. The process is governed by a mean function and a covariance (kernel) function. The mean function $m(x)$ is your best guess for the function shape before seeing any data.

Setup : We define :

— A set of input points $x \in \mathbb{R}^{100}$, evenly spaced in the interval [-5, 5];

— For our example, we define a mean function $m(x) = \frac{1}{4}x^2$, representing the expected value of the function at each point;

— A kernel function given by the squared exponential (RBF) form :

$k(x, x') = \exp\left(-\frac{1}{2}(x - x')^2\right)$, which defines the entries of the covariance matrix $\Sigma \in \mathbb{R}^{100\times100}$ by evaluating the kernel at all pairs of input points.

Sampling : We draw 3 independent samples $f^{(1)}, f^{(2)}, f^{(3)} \in \mathbb{R}^{100}$ from the multivariate normal distribution : $f^{(i)} \sim \mathcal{N}(\mu, \Sigma)$, where each sample represents one possible realization of a function under the GP prior.

Uncertainty Visualization : The 95% confidence interval around the mean is visualized by shading the region :

$$m(x) \pm 1.96\sqrt{\mathrm{diag}(\Sigma)},$$

where $\sqrt{\mathrm{diag}(\Sigma)}$ represents the standard deviation at each input point.

Interpretation : This example highlights a key feature of GPs : instead of fitting a single curve to data, GPs define a distribution over functions. Each sampled curve represents a plausible function consistent with the prior defined by m and Σ, even before observing any data. Figure 2.8 illustrates several such realizations : $f^{(1)}$ (Sample 1 in blue), $f^{(2)}$ (Sample 2 in green) and $f^{(3)}$ (Sample 3 in red).

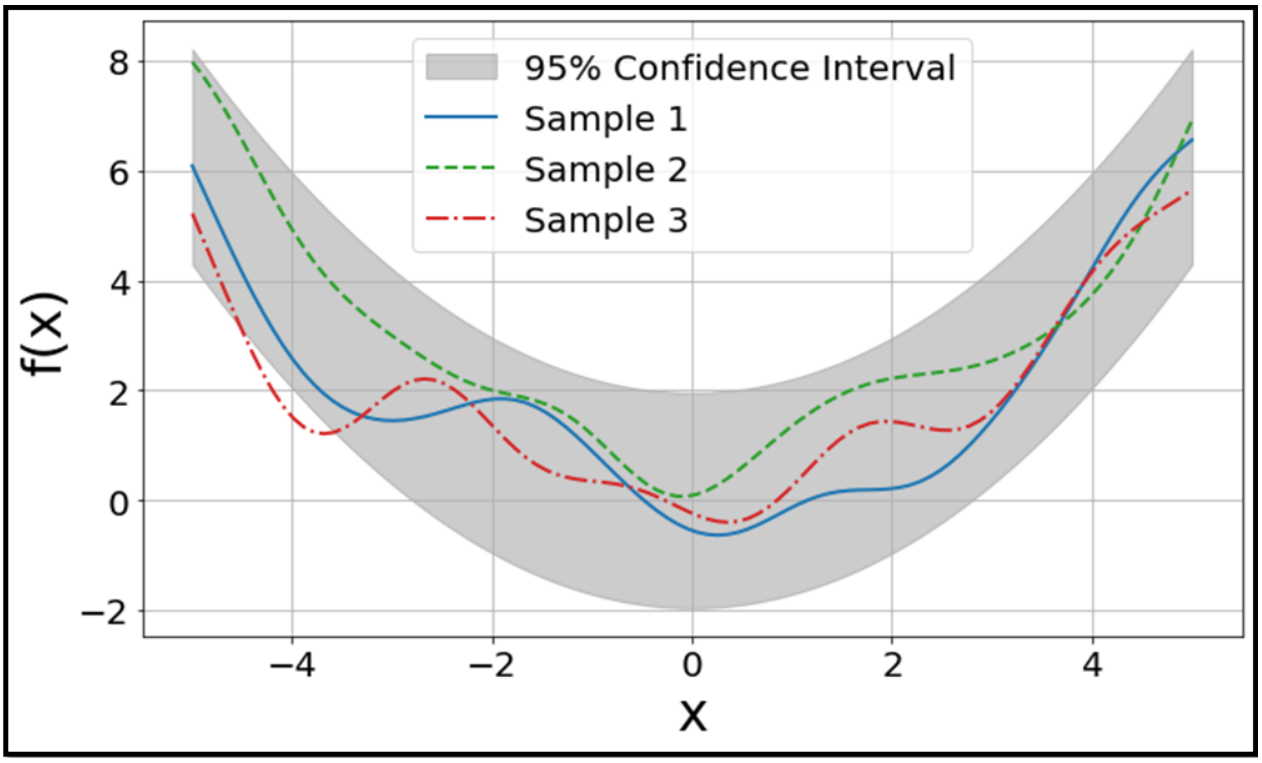

FIGURE 2.8 : Gaussian Process samples with confidence interval.

Three sample functions drawn from a Gaussian Process with mean function $m(x) = \frac{1}{4}x^2$ and squared exponential kernel $k(x, x') = \exp\left(-\frac{1}{2}(x - x')^2\right)$. Colored curves represent distinct realizations. The shaded area is the 95% confidence interval $m(x) \pm 1.96\sqrt{\mathrm{Var}[f(x)]}$.

Here, we considered GPs as priors over functions, defined solely by a mean function and a covariance kernel. However, once data is observed, the GP framework allows us to compute a posterior distribution over functions that are consistent with both the prior and the data. This posterior forms the basis of GPR, enabling both prediction and quantification of uncertainty at unseen inputs.

Gaussian Process posterior

In the previous section, we introduced GPs as distributions over functions, defined by a mean function and a covariance kernel. When no data is available, these priors capture our initial beliefs

about function smoothness, amplitude, and variability. However, once observations are made, we can update these beliefs. This process yields the GP posterior, which forms the basis of GPR.

Given training data $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$, where each y_i is a noisy observation of the underlying function $f(x)$, we model the observations as :

Equation 2.3 : Noisy observation model for Gaussian Process regression

$$y_i = f(x_i) + \varepsilon_i, \qquad \varepsilon_i \sim \mathcal{N}\left(0, \sigma_n^2\right) \tag{2.3}$$

The GP posterior provides a distribution over function values at any new test point x_*, which remains Gaussian. Its predictive mean and variance are given by :

Equation 2.4 : Posterior predictive mean and variance in Gaussian Process regression

$$\mu(x_*) = k(x_*, X)\left[K(X, X) + \sigma_n^2 I\right]^{-1}\mathbf{y}, \tag{2.4}$$

$$\sigma^2(x_*) = k(x_*, x_*) - k(x_*, X)\left[K(X, X) + \sigma_n^2 I\right]^{-1}k(X, x_*). \tag{2.5}$$

where :

— σ_n^2 is the noise variance,

— $K(X, X)$ is the $n \times n$ covariance matrix between training inputs,

— $k(x_*, X)$ is the $1 \times n$ vector of covariances between the test input x_* and all training inputs,

— $\mathbf{y}$ is the vector of observed targets.

The predicted mean $\mu(x_*)$ is a linear combination of the observed outputs, weighted by their similarity to x_* as encoded by the kernel. $\sigma^2(x_*)$ is the predictive variance, which measures the uncertainty in the prediction. This uncertainty is naturally lower near regions with training data and increases in regions far from observations.

Overall, GPR provides not only a flexible model for non-linear regression but also a principled way to quantify uncertainty.

Advantages of Gaussian Process Regression :

— **Uncertainty quantification** : GPR provides variance estimates along with predictions.

— **Non-parametric** : The model complexity grows with the data.

— **Data efficiency** : Performs well with small datasets.

— **Interpretability** : The kernel encodes meaningful structural assumptions.

Drawbacks of Gaussian Process Regression :

— **Computational cost :** Involves inverting an $n \times n$ matrix, which becomes expensive for large datasets ($\mathcal{O}(n^3)$).
— **Scalability :** Not well suited for datasets with tens or hundreds of thousands of points without approximation methods.
— **Kernel selection :** Choosing and tuning the right kernel requires domain knowledge and experience.

Therefore, GPR is used in :

— BO for experimental design and tuning,
— Surrogate modeling of expensive simulations,
— Geostatistics,
— Uncertainty-aware learning in physics and engineering.

2.5.6 Bayesian Optimization with Gaussian Processes

BO is a probabilistic global optimization framework designed for optimizing objective functions that are expensive to evaluate, noisy, or lack a closed-form expression. The core principle involves modeling the unknown objective function $f(x)$ with a surrogate model, typically a GP, that provides a posterior predictive mean and variance.

At each iteration, BO uses this surrogate to define an acquisition function that guides the selection of the next evaluation point, balancing two competing objectives :

— **Exploration** : Sampling points in regions of high uncertainty to gain knowledge of the function.
— **Exploitation** : Sampling where the predicted mean is favorable to refine known good solutions.

Bayesian Optimization Procedure

1. **Surrogate modeling :** Initialize a Gaussian Process at a time t using the observed dataset $\mathcal{D}_t = \{(x_i, f(x_i))\}_{i=1}^{t}$, yielding a posterior predictive distribution :

Equation 2.5 : Posterior GP surrogate model used in BO

$$f(x) \sim \mathcal{GP}\big(\mu_t(x),\, \sigma_t^2(x)\big) \tag{2.6}$$

2. **Acquisition function maximization :** Define an acquisition function $\alpha_t(x)$ and select the next evaluation point as :

Equation 2.6 : Next evaluation point in BO

$$x_{t+1} = \arg\max_x \alpha_t(x) \tag{2.7}$$

3. **Function Evaluation :** Query the true objective function $f(x_{t+1})$ at the new point.
4. **Model Update :** Augment the dataset and retrain the GP.
5. **Repeat** until convergence or budget exhaustion.

The BO framework is given in Algorithm 1, where we iterate between steps 2 and 5 until the evaluation is over (with $\mathcal{D}_{1:n} = \{(\mathbf{x}_1, y_1)\}, \{(\mathbf{x}_2, y_2)\}, \{(\mathbf{x}_3, y_3)\}, ..., \{(\mathbf{x}_n, y_n)\}$) :

Algorithm 1 Bayesian Optimization

1: **for** $n = 1, 2, 3, \ldots$ **do**
2: Select new $\mathbf{x}_{n+1}$ by optimizing the acquisition function α :

$$\mathbf{x}_{n+1} = \arg\max_{\mathbf{x}\in\mathcal{X}} \alpha(\mathbf{x}; \mathcal{D}_{1:n})$$

3: Query the objective function to obtain y_{n+1}
4: Augment data $\mathcal{D}_{1:(n+1)} = \mathcal{D}_{1:n} \cup \{(\mathbf{x}_{n+1}, y_{n+1})\}$
5: Update the GP model
6: **end for**

Common Acquisition Functions

Several acquisition functions have been proposed to formalize the exploration-exploitation trade-off :

— **Expected Improvement (EI) :**

$$\alpha_{\mathsf{EI}}(x) = \mathbb{E}[\max(f(x) - f(x_*), 0)] = (\mu(x) - f(x_*))\Phi(Z) + \sigma(x)\phi(Z)$$

where the x_* refers to the best point found so far, $\mathbb{E}[\cdot]$ denotes the mathematical expectation, $Z = \frac{\mu(x)-f(x_*)}{\sigma(x)}$, Φ is the standard normal cumulative distribution function (CDF), and ϕ is the probability density function (PDF).

— **Probability of Improvement (PI) :**

$$\alpha_{\mathsf{PI}}(x) = \Phi\left(\frac{\mu(x) - f(x_*) - \xi}{\sigma(x)}\right)$$

where ξ is a small positive parameter to encourage exploration.

— **Upper Confidence Bound (UCB) :**

$$\alpha_{\text{UCB}}(x) = \mu(x) + \kappa\sigma(x)$$

where $\kappa > 0$ balances exploration (high κ) and exploitation (low κ).

Illustrative Example of Bayesian Optimization (Figure 2.9)

To clarify how BO unfolds in practice, we illustrate a concrete example corresponding to Figure 2.9. The goal is to maximize the one-dimensional objective function : $f(x) = \sin(3x) + 0.3\cos(5x)$, defined over the interval $x \in [0, 7]$. This function is deterministic, but assumed to be expensive to evaluate. To simulate a real-world optimization task, we begin with only two randomly chosen initial evaluations of $f(x)$, which form the initial dataset $\mathcal{D}_2 = \{(x_1, y_1), (x_2, y_2)\}$.

A GP surrogate model is trained on this initial dataset, yielding a posterior mean $\mu(x)$ and uncertainty $\sigma(x)$ over the domain. Using these quantities, we construct the UCB acquisition function : $\text{UCB}(x) = \mu(x) + \kappa\sigma(x)$, which balances exploitation of promising regions (high $\mu(x)$) and exploration of uncertain areas (high $\sigma(x)$). The next sampling location x_3 is chosen by maximizing UCB.

This procedure is repeated iteratively : at each iteration, a new point x_{t+1} is selected, the true function is evaluated at that point, and the dataset is updated with $(x_{t+1}, f(x_{t+1}))$. The GP model is retrained on the new dataset, and the acquisition function is recomputed accordingly. Figure 2.9 corresponds to iteration 12 of this process; UCB values at iteration 12 are represented in Table 2.4 where κ is kept constant during the whole process. The figure shows :

— The true target function $f(x)$ (black dashed curve),
— The GP posterior mean (blue curve) and its 95% confidence interval (blue shaded region),
— The acquisition function (green dashed line),
— The past observations (red dots),
— The next sampling point (black circle), which corresponds to the maximum of the UCB.

TABLE 2.4 : UCB values at iteration 12

Iteration	$\mu(x_{12})$	$\sigma(x_{12})$	κ	$\text{UCB}(x_{12})$
12	1.28	0.033	2	1.34

This example concretely illustrates how BO dynamically balances exploration and exploitation to efficiently locate the maximum of an unknown function.

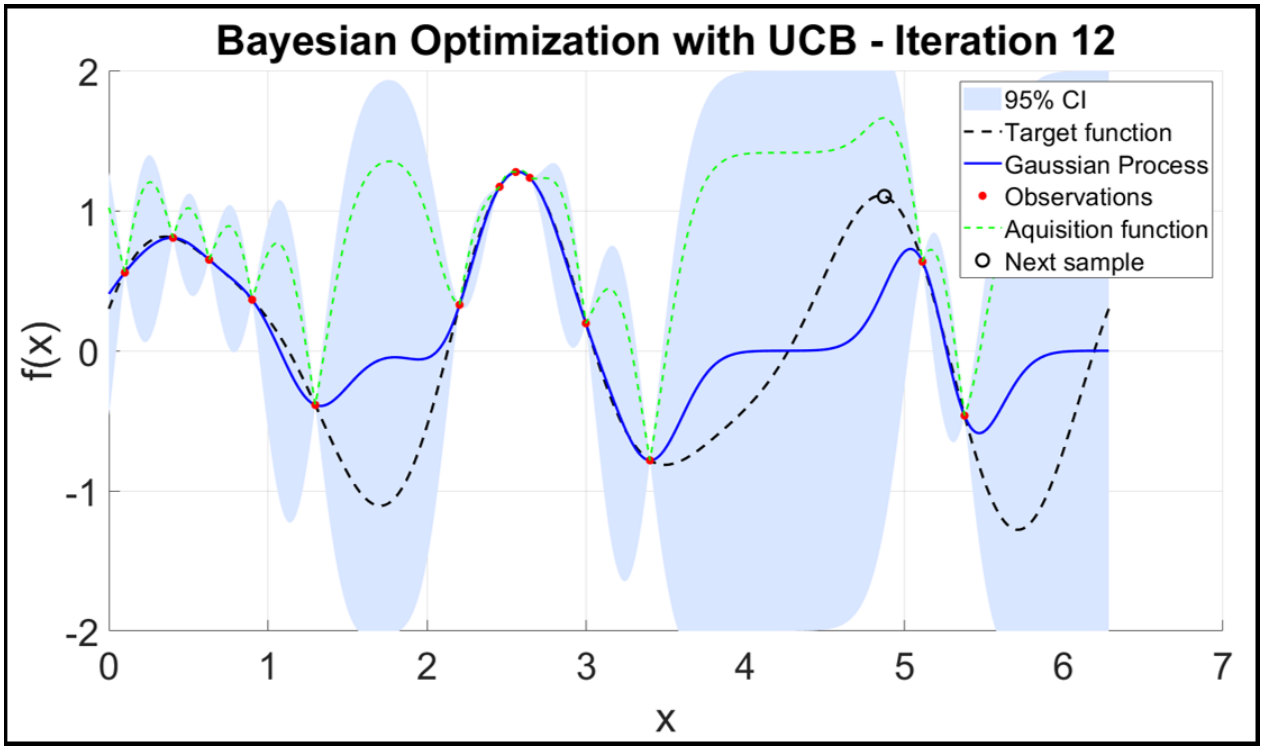

FIGURE 2.9 : Example of GP modeling and UCB-guided sampling at iteration 12.

The true objective function $f(x)$ is shown as a dashed black curve. The GP surrogate model is plotted in blue, with a shaded region representing the 95% confidence interval. Red dots indicate past observations, while the green dashed line represents the UCB acquisition function. The next sampling point, selected as the maximizer of the UCB, is marked with a black circle.

Scientific Applications

BO is well-suited to scenarios where objective evaluations are costly or limited in data points. It has been applied extensively to :

— Hyperparameter tuning in ML,

— Surrogate-assisted design in engineering,

— Material discovery and drug design,

— Experimental optimization in physics and biology.

In high-power laser experiments, where each shot is costly and setup changes are time-consuming, BO is a sample-efficient strategy for identifying optimal configurations, such as temporal pulse shape or beam wavefront, while minimizing the number of required iterations. Its ability to incorporate uncertainty estimates from GPR makes it particularly well-suited for closed-loop optimization in complex, noisy environments.

BO, in turn, leverages these surrogate models to optimize black-box objective functions that are expensive to evaluate, noisy, and potentially non-differentiable. It is particularly effective in low-dimensional spaces (typically ≤ 20 parameters) (Jones et al., 1998; Mockus, 2005; Brochu et al., 2010; Frazier, 2018), making it ideal for experimental plasma physics applications where rapid convergence is essential.

Recent studies have demonstrated the effectiveness of GPR in plasma physics. Döpp et al. (2023) compared GPR with other interpolation techniques such as nearest-neighbor and cubic splines, showing that GPR provides superior predictive accuracy with well-defined uncertainty estimates. In another study, McClarren et al. (2011) used a physics-informed emulator based on GPR and Bayesian Multivariate Adaptive Regression Splines (BMARS) to model uncertainty in laser-driven radiating shock simulations. While BMARS effectively handled discontinuities, GPR excelled in uncertainty quantification, making it a valuable tool for sensitivity analysis and optimization.

Numerical approaches leveraging Bayesian methods have emerged as powerful tools for optimizing laser–plasma interactions and predicting proton beam characteristics. BO is particularly effective in exploring large, high-dimensional parameter spaces while minimizing computational cost. Ferran Pousa et al. (2023) demonstrated a BO framework that combined high-fidelity PIC simulations (*FBPIC*) with reduced-model simulations (*Wake-T*), achieving a tenfold speedup in optimization. Similarly, Dolier et al. (2022) applied BO with GPR in the code *BISHOP*, optimizing laser-driven proton acceleration by efficiently tuning laser energy, pulse duration, and target properties to maximize proton energy reaching 220 MeV. Further, Takagi et al. (2021) introduced a statistical approach using Bayesian inference to predict the maximum proton energy achievable through TNSA. By analyzing experimental parameters alongside hot electron temperature and density from corresponding PIC simulations, the study derives a multivariate scaling law that accurately forecasts proton energies under varying experimental conditions. This methodology not

only demonstrates the effectiveness of Bayesian inference in modeling complex laser-plasma interactions but also provides valuable insights for designing experiments aimed at achieving proton energies exceeding 100 MeV.

In experimental applications, Loughran et al. (2023) showcased the first real-time BO of a laser-driven ion source, achieving performance comparable to manual tuning but with only 57% of the laser energy consumption. Other studies have applied BO to betatron radiation tuning in LWFA, improving X-ray photon yields within just a few iterations. By tuning plasma density and laser focal position, Ye et al. (2022) applied BO to efficiently optimize betatron radiation from LWFA. They used GPR to rapidly identify optimal conditions, doubling the X-ray photon number within 10 iterations, compared to traditional grid searches. A PIC simulation further explained the observed radiation enhancement, showing how plasma density influences electron injection and oscillation amplitude. A key advantage of BO in accelerator physics is its adaptability to multi-objective optimization. Irshad et al. (2023) developed a Bayesian framework that balances competing parameters, constructing a Pareto front (Tušar et al., 2014) for optimizing electron beam energy and charge while minimizing computational costs. Furthermore, Shalloo et al. (2020) automated the tuning of six key laser wakefield accelerator parameters : pulse energy, chirp, pulse duration, focal position, plasma density, and target length. This optimization led to an 80% increase in electron beam charge while adjusting pulse duration by just 1%.

The increasing role of BO in plasma physics extends to safe and efficient optimization techniques. Jiang et al. (2025) implemented parallel BO for compact FELs driven by laser wakefield accelerators in the soft X-ray regime. Using GPR as a surrogate model, the method efficiently tunes FEL beamline parameters to maximize radiation energy. Starting from random parameter settings, the optimization process increased FEL energy from 10 nJ to 25 μJ within tens of iterations.

GPR and BO offer powerful tools for modeling, optimization, and uncertainty quantification in the complex, high-dimensional parameter spaces of laser–plasma physics. GPR provides both predictive accuracy and confidence intervals, enabling the construction of surrogate models that capture nonlinear dependencies and reduce reliance on computationally expensive PIC simulations. BO leverages GPR to optimize experimental and simulation parameters with minimal evaluations, significantly lowering the cost of parameter sweeps : an advantage when trials are expensive or high-fidelity simulations are time-consuming. Limitations include GPR's poor scalability with large datasets ($O(n^3)$ complexity) and sensitivity to kernel choice, as well as BO's reduced efficiency in high-dimensional spaces (>20 variables) without dimensionality reduction or physics-informed priors. Standard BO can also suffer from limited interpretability, motivating hybrid approaches that integrate domain knowledge. Despite these challenges, combining GPR and BO with multi-fidelity models, adaptive experimental design, and real-time optimization frameworks has shown promise in optimizing plasma-based accelerators. As computational power and dataset sizes grow, these methods are poised to play a central role in improving predictive plasma models, enhancing experimental reproducibility, and accelerating discovery in high-energy-density physics.

2.5.7 Reinforcement learning

Reinforcement Learning (RL) is a branch of ML focused on how agents can learn to make sequences of decisions by interacting with an environment. Unlike Supervised Learning, where learning is guided by known input-output pairs, RL involves learning through trial and error, guided by feedback in the form of rewards or penalties. This approach is particularly well-suited to problems where the correct actions are not immediately obvious and must be discovered over time. Foundational work in this field is well-documented in Sutton et al. (1998).

The basic RL setup includes an agent that perceives its environment, selects actions, and receives feedback. At each step, the agent observes the current situation, chooses an action, and obtains a reward from the environment. This action also changes the environment's state, prompting the next decision. Over time, the agent uses this feedback to develop strategies that maximize cumulative reward. The environment and agent together form a feedback loop that makes RL dynamic and adaptable represented in Figure 2.10.

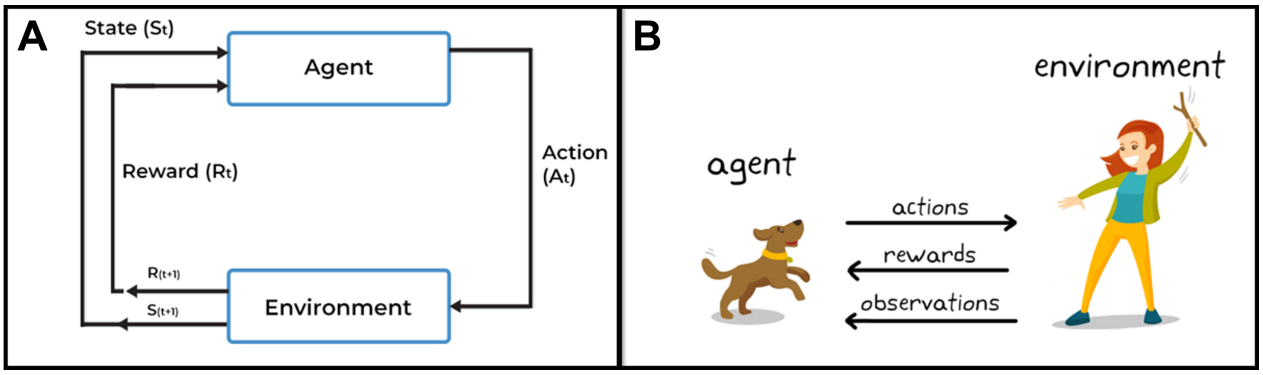

FIGURE 2.10 : Reinforcement Learning interaction loop and analogy.

(A) RL workflow which shows the formal agent-environment interaction loop where the agent receives a state S_t, selects an action A_t, and obtains a reward R_t, influencing the next state S_{t+1}. (B) The goal of learning in this case is to train the dog (agent) to complete a task within an environment, which includes the surroundings of the dog as well as the trainer, highlighting the roles of actions, observations, and rewards in RL. These figures have been adapted from Spiceworks (2024) and The MathWorks Inc. (2024), with the publishers' permission.

RL algorithms are broadly divided into model-free and model-based approaches. Model-free algorithms like Q-learning and policy gradient methods learn directly from experience without constructing an explicit model of the environment's dynamics (Watkins et al., 1992; Schulman et al., 2017). In contrast, model-based methods attempt to learn how the environment behaves, enabling the agent to simulate and plan future actions (Ha et al., 2018). A hybrid strategy known as actor-critic combines policy learning (actor) with value estimation (critic), improving training efficiency and stability.

A central challenge in RL is the exploration-exploitation dilemma. To maximize rewards, the agent must exploit known strategies that perform well. However, to discover better strategies, it must also

explore unknown actions. This balance is maintained using techniques such as random action selection, uncertainty-driven sampling, or entropy regularization. Effective exploration strategies are crucial for success, particularly in large or unfamiliar environments.

RL has delivered groundbreaking results in a variety of domains. In games, RL systems such as *AlphaGo* and *AlphaZero* have achieved superhuman performance in Go, Chess, and Shogi by combining RL with deep neural networks and self-play (Silver et al., 2016, 2018). In robotics, RL allows machines to learn physical tasks like walking or object manipulation in dynamic settings (Kober et al., 2013). In healthcare, RL has been used to design personalized treatment strategies and optimize chemical reactions in drug discovery (Zhou et al., 2017). Autonomous vehicles also use RL to learn complex navigation and control tasks.

Despite its versatility, RL faces important challenges. One major limitation is sample inefficiency. Many algorithms require vast amounts of experience to learn effectively. Generalizing learned behaviors to novel situations remains difficult, especially in high-dimensional spaces. Furthermore, RL decisions are often difficult to interpret, raising concerns in safety-critical applications. Techniques like Deep Q-Networks (DQN) and Proximal Policy Optimization (PPO) have helped RL scale to more complex and high-dimensional problems (Mnih et al., 2015).

There, RL offers a powerful and flexible framework for adaptive decision-making. By learning from interaction and focusing on long-term outcomes, RL is uniquely positioned to solve problems where dynamic, context-sensitive behavior is essential. As research continues to address current limitations and improve sample efficiency and safety, RL has emerged as a promising alternative, offering model-free optimization capabilities that enable real-time, adaptive control. Unlike conventional methods, RL continuously learns and refines control policies based on direct interaction with the system.

Deep RL (DRL) extends traditional RL by incorporating deep neural networks, enabling it to handle high-dimensional and complex state spaces that are unmanageable for tabular or linear methods. Both aim to teach an agent how to act based on rewards from its environment. The key difference lies in how they process information. RL works with simple, low-dimensional data using rule-based tables, while DRL uses deep neural networks to handle complex, high-dimensional inputs like images or sensor data. For example (Table 2.5), imagine teaching a robot to ride a bike :

— In RL, the robot uses a few basic rules like “if tilting left, turn right.”
— In DRL, the robot learns directly from a video feed using a neural network to decide what to do.

One of the key applications of DRL is the optimization of the temporal profile of high-power laser pulses. Traditional pulse shaping methods rely on iterative simulations, which can be computationally demanding and fail to account for transient effects. In contrast, DRL enables a data-driven approach to autonomously adjust dispersion coefficients, effectively minimizing pulse duration and maximizing intensity. For example, Capuano et al. (2025) developed a DRL framework that successfully optimized laser pulse shaping at the *ELI Beamlines* facility, achieving near transform-

TABLE 2.5 : Simple comparison between RL and DRL.

Aspect	Reinforcement Learning (RL)	Deep Reinforcement Learning (DRL)
Input type	Simple, low-dimensional data	Raw, complex inputs (e.g., images, sensors)
Model type	Lookup table or basic function	Deep neural network
Best for	Simple environments with few states	Complex, high-dimensional environments
Example	"If tilting left, turn right" rule	Learns from camera feed to keep balance

limited pulses — i.e., pulses compressed close to the minimum duration allowed by their spectral bandwidth — with a full width at half maximum (FWHM) of 1.6 ps. The RL agent, trained using a high-power laser simulator, was capable of adjusting spectral phase coefficients in real time to optimize intensity while maintaining pulse stability. In another study, *TempoRL* (Capuano et al., 2023), explored a similar approach for adaptive dispersion control using DRL, demonstrating that model-free learning can effectively tune nonlinear temporal dynamics without requiring explicit physical models. Model-free learning is like learning by trial and error. The agent improves its behavior just by interacting with the environment, without needing to understand how the environment works.

In nonlinear optics, RL has been employed for self-tuning and control of mode-locked fiber lasers, which are critical for ultrafast laser applications. Traditional tuning methods often require direct parameter sweeps, which can be inefficient. Sun et al. (2020) integrated DRL to autonomously adjust waveplate and polarizer settings, achieving robust mode-locking even under fluctuating system conditions. A comparative study by Pan et al. (2025) evaluated three RL algorithms - Double Deep Q Network (DDQN), Twin Delayed Deep Deterministic Policy Gradient (TD3), and Soft Actor-Critic (SAC) - for optimizing ultrafast mode-locked lasers. Their results demonstrated that :

— DDQN effectively learned control policies but required longer training.

— TD3 converged more quickly and efficiently handled continuous action spaces.

— SAC outperformed both, providing the highest stability and robustness, even under disturbances.

These findings suggest that RL, particularly SAC-based methods, offers a powerful framework for adaptive photonic system control, enhancing both performance and reliability.

RL has also been successfully applied to conventional particle accelerator optimization, where beam alignment and trajectory control are crucial. At *FERMI FEL*, RL techniques such as Q-learning and Natural Policy Gradient algorithms have been employed to align seed lasers and maintain beam stability by compensating for thermal drifts and perturbations in real-time (Bruchon et al., 2020). Similarly, *CERN* has implemented model-free RL algorithms for trajectory steering in electron lines and linear accelerators, demonstrating the ability of RL to adapt to time-varying system dynamics and handle multiple degrees of freedom (Kain et al., 2020). These applications

highlight RL's effectiveness in automating traditionally manual alignment processes, significantly improving efficiency and beam quality.

Deploying RL in high-power laser laboratories demands careful safeguards. Policies must remain within hardware and physics constraints, requirements that can be enforced through safe-RL formulations (algorithms specifically designed to avoid unsafe actions during learning), rule-based guards (predefined safety limits or thresholds coded into the system), or hybrid schemes (approaches that combine RL with physics-based models or conventional control logic to guide learning within safe operational boundaries). Because each laser shot is expensive, data efficiency is paramount : model-based RL, transfer learning, or physics-informed priors can extract maximum value from limited experiments. Smooth integration with existing control software is also essential, often achieved by coupling fast, interpretable physics models with RL agents that refine those baseline solutions. With these measures in place, RL becomes a transformative tool for optimizing laser-based infrastructures, elevating the performance and reproducibility of high-energy-density physics experiments.

Future improvements in RL should aim to :

— Make learning faster and more efficient by reusing knowledge from past experiments (using methods like meta-learning (Schweighofer et al., 2003) and transfer learning (Weiss et al., 2016).
— Include basic physics rules in the learning process to make results more stable and easier to understand.
— Use safety-focused RL methods to keep experiments within safe and reliable limits.

With continued advancements, RL-driven control systems have the potential to fully automate complex experimental setups, accelerating discoveries in laser-plasma physics and high-intensity lasers.

2.6 Conclusion and perspectives on Machine Learning in laser-driven ion acceleration

2.6.1 Impact on Experimental Efficiency

The integration of ML techniques into laser-plasma acceleration has improved experimental throughput and data quality in several ways :

1. **Reduction in experiment time :**
   — With BO, optimization is achieved in fewer iterations than traditional parameter sweeps especially in multi-dimensional parameter spaces.

2. **Enhanced data interpretation and uncertainty quantification :**
   - GPR models provide probabilistic confidence intervals, quantifying and taking into account experimental uncertainty.
   - Neural network-based synthetic diagnostics allow real-time feedback, preventing wasted laser shots on suboptimal configurations.
3. **Adaptive control for complex systems :**
   - RL-based control systems can automatically fine-tune experimental parameters in real time, helping maintain consistent beam properties even under changing conditions.
   - Combining adaptive optics with ML allows for real-time correction of wavefront distortions (Loughran et al., 2023), improving laser pointing stability (Amodio et al., 2025) and enhancing ion acceleration efficiency (Catrix et al., 2025).

2.6.2 Insights from Machine Learning-driven optimization

Beyond optimizing existing techniques, ML has shown insights into laser-plasma interactions :

1. **Identification of hidden correlations :**
   - Deep learning models discovered that pre-plasma gradient length plays a more significant role in proton energy than previously understood (Djordjević et al., 2021a).
   - SVR models revealed that the optimal thickness for TNSA targets varies based on pulse contrast and electron temperature, a dependency previously overlooked (Desai et al., 2024).
2. **New scaling laws for ion acceleration :**
   - Bayesian inference models led to revised scaling laws for TNSA, demonstrating that multi-variable regression models provide more accurate predictions than traditional power-law approximations (Takagi et al., 2021).
3. **Optimization of multi-objective trade-offs :**
   - BO-enabled optimization of laser-plasma interactions led to Pareto-optimal solutions balancing proton energy, beam divergence, and pulse duration (Irshad et al., 2023).

ML techniques have transformed LDIA by improving predictive modeling, experimental efficiency, and real-time adaptability. While neural networks excel at pattern recognition, RFs deliver noise-resilient nonlinear surrogate modeling along with quantitative feature importance, SVR offers computationally efficient regression in low-data regimes, GPR enhances uncertainty quantification, and RL enables adaptive control. Table 2.6 summarizes these methods' roles and trade-offs.

Looking forward, further progress in LDIA is likely to come from several emerging directions :

— **Hybrid ML approaches :** Combining techniques such as GPR and RL can enhance data efficiency while enabling adaptive control strategies that respect experimental constraints.
— **Physics-informed neural networks (PINNs) :** Embedding known physical laws into neural networks improves interpretability and reduces the amount of data needed for training.
— **Experiment-driven ML frameworks :** Integrating synthetic diagnostics with closed-loop optimization creates robust systems capable of adapting in real-time to changing conditions.
— **Real-time feedback mechanisms :** Advanced control systems using RL or BO can fine-tune laser parameters on the fly, improving beam quality and reproducibility.
— **High-performance computing (HPC)-integrated simulations :** Coupling ML with high-performance computing accelerates parameter scans, surrogate model training, and uncertainty quantification in complex simulations.

These strategies aim to bridge the gap between modeling and experimentation, pushing laser-plasma accelerators toward full automation, enhanced beam quality, and faster experimental cycles. Continued interdisciplinary collaboration between plasma physicists, computer scientists, and ML experts will be essential to drive these advances forward.

TABLE 2.6 : Comparison of ML techniques for optimization.

ML method	Most Efficient Domain	Key Strengths	Key Weaknesses
Neural Networks (NN)	Predictive modeling of ion beam properties; synthetic diagnostics.	Handles complex, high-dimensional parameter spaces; generalization for unseen data.	Requires large datasets; computationally expensive training; black-box nature limits interpretability.
Support Vector Regression (SVR)	Regression tasks with limited data; low-dimensional parameter space optimization.	Computational efficiency; works well in low-data environments; avoids overfitting.	Poor scalability for large datasets; limited extrapolation ability beyond training range.
Random Forest Regression (RF)	Surrogate modeling and sensitivity analysis for sparse, noisy experimental datasets.	Robust to noise and shot-to-shot fluctuations; good generalization with limited data; provides feature importance for identifying key laser/target parameters.	Limited extrapolation capability outside training domain; ensemble averaging reduces interpretability compared to single-tree models.
Gaussian Process Regression (GPR)	Uncertainty quantification; surrogate modeling for simulations; probabilistic modeling.	Provides confidence intervals; useful for experimental planning; surrogate modeling.	Computationally expensive for large datasets ($O(n^3)$ complexity); limited scalability.
Bayesian Optimization (BO)	Global optimization of laser and target parameters; tuning experimental setups.	Reduces experimental iterations; balances exploration-exploitation in parameter tuning.	Less effective in high-dimensional parameter spaces; requires good surrogate models.
Reinforcement Learning (RL)	Adaptive real-time control; high-dimensional optimization with dynamic constraints.	Learns from direct interaction; continuously improves experimental conditions dynamically.	Sample inefficient (requires many iterations); safety constraints for experimental use.

Each method has specific strengths and limitations, making them suitable for different applications within experimental and computational laser-driven ion acceleration.

3 STABLE HIGH REPETITION-RATE LASER-DRIVEN PROTON BEAM PRODUCTION FOR MULTIDISCIPLINARY APPLICATIONS

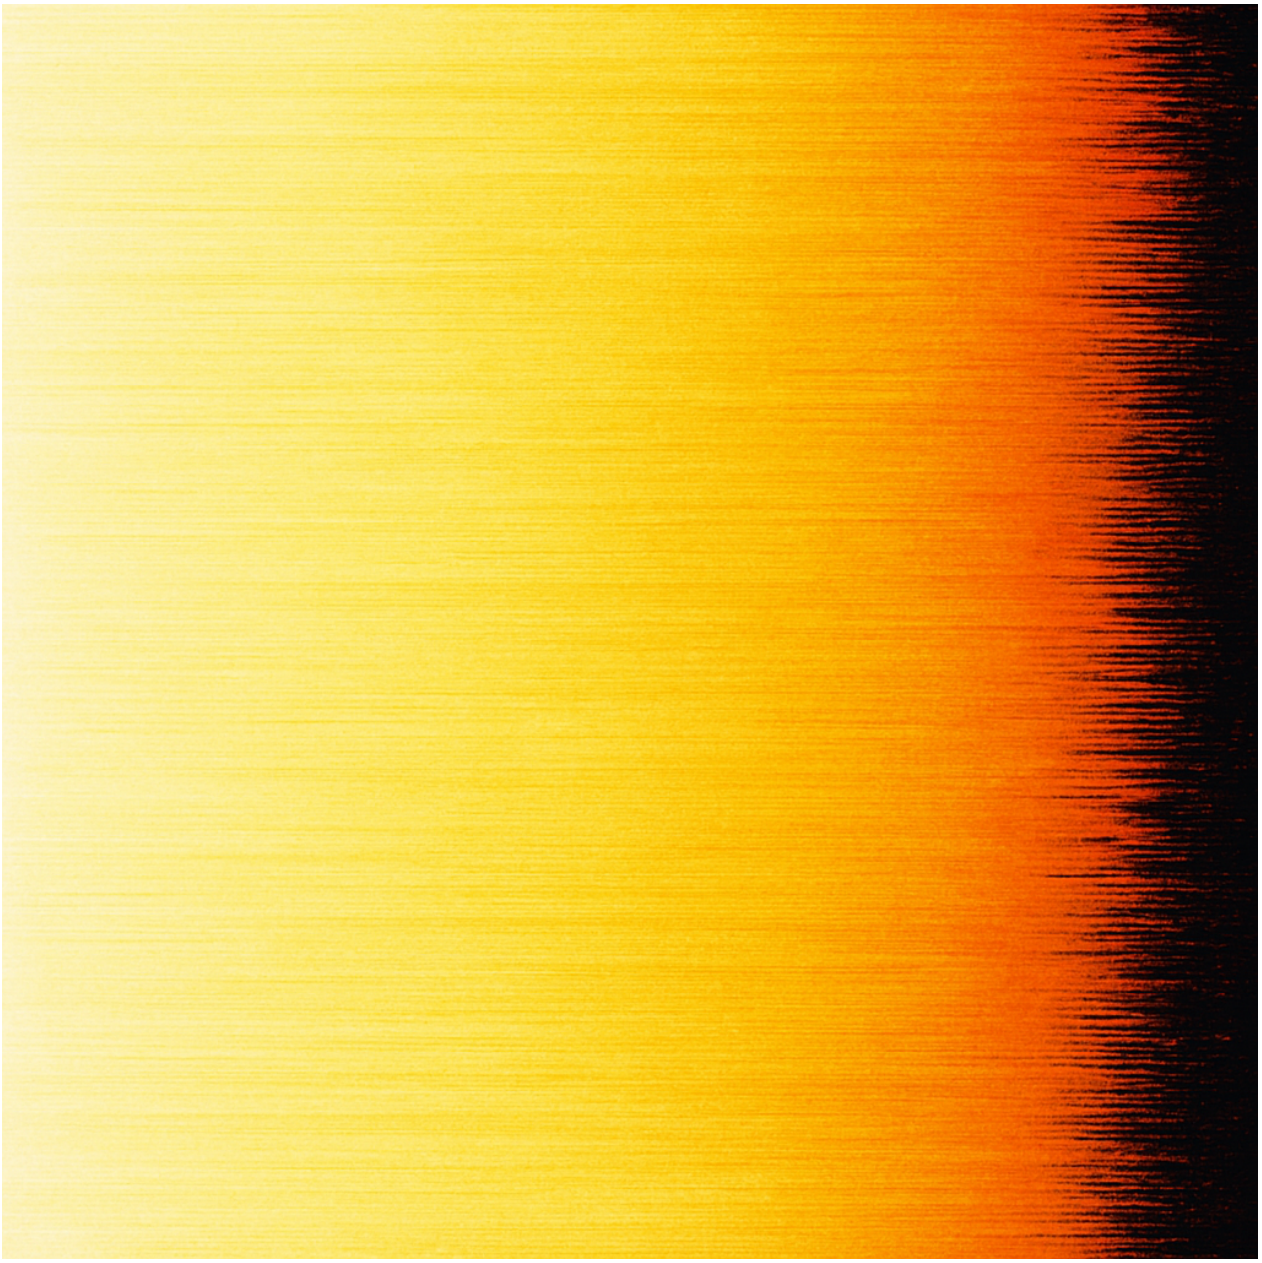

Stable sequence of consecutive laser-driven proton beam spectra. Original image created by the author (Elias Catrix). No external permission is required; use authorized by the author.

Foreword — This chapter builds upon and expands the findings published in the journal *Review of Scientific Instruments*, in the article titled "Stable high-repetition-rate laser-driven proton beam production for multidisciplinary applications on the Advanced Laser Light Source ion beamline" (Catrix et al., 2023). It further develops the underlying concepts and incorporates additional insights that extend beyond the original publication.

Laser-driven ion accelerators are increasingly regarded as compact and versatile tools for a wide range of applications. The advent of multi-hundred-TW laser systems has enabled the routine generation of relativistic peak intensities above 10^{18} W.cm^{-2}, producing ion beams in the multi-MeV energy range. Yet, challenges remain in achieving the stability, reproducibility, and beam quality required for practical deployment, especially in contexts demanding narrow spectral control or high ion flux.

This chapter describes a high-repetition-rate (0.625 Hz) laser-based ion acceleration platform implemented at the 150 TW Ti : Sapphire beamline of the *Advanced Laser Light Source* (*ALLS*). The system delivers laser pulses with an RMS energy stability of 2.5%, focused to intensities between 10^{19} and 10^{20} W.cm^{-2}. Proton acceleration in the Target Normal Sheath Acceleration regime was studied using solid foils of various compositions and thicknesses. Through systematic multi-shot series — up to 400 laser shots per pumping cycle — we quantify fluctuations in maximum proton energy, proton yield, energy-integrated proton yield, and laser-to-proton energy conversion efficiency across different target configurations. These high-statistics datasets provide critical insights into the operational stability of laser–plasma accelerators, enabling their use in high-throughput applications where reproducibility is paramount.

Building on this foundation, the chapter explores one such application : laser-driven Particle-Induced X-ray Emission (PIXE), where enhanced beam repeatability — supported by the careful selection of target material and thickness — directly translates into improved accuracy and reliability in trace element analysis.

3.1 Introduction

The practical deployment of laser-driven ion accelerators is hindered by the inherently pulsed and fluctuating nature of laser–plasma ion sources. Among the various acceleration mechanisms, Target Normal Sheath Acceleration (TNSA) remains the most widely studied and reliably reproduced, typically using micrometer- or nanometer-thick solid targets to generate multi-MeV ion beams. Achieving high laser-to-ion energy conversion efficiency in this regime requires stringent conditions, particularly in laser focusing and target alignment. To maximize on-target intensity, solid targets must be positioned within the laser's focal plane with sub-Rayleigh length precision, typically on the order of $z_R = \frac{\pi w_0^2}{\lambda} = \frac{\pi\left(\frac{w_{\mathrm{FWHM}}}{2\sqrt{\ln 2}}\right)^2}{800\,\mathrm{nm}} \approx 70.8\,\mu$m at the *Advanced Laser Light Source* (*ALLS*), where $w_{\mathrm{FWHM}} \sim 5$ μm is the focal spot size at the Full Width at Half Maximum (FWHM),

placing strict constraints on shot-to-shot alignment. Compounding this challenge, each target is destroyed after a single laser shot, necessitating a rapid and automated delivery system to sustain high-repetition-rate operation. Without such automation, ion flux is severely limited by the time-consuming process of manual replacement, undermining the practicality of laser-driven ion acceleration in high-throughput experimental or applied environments.

To address these constraints, the laser-plasma community has developed a range of strategies to enable higher repetition rates. Tape targets (Noaman ul Haq et al., 2017; Noaman ul Haq et al., 2018; Dover et al., 2020b) offer a cost-effective and simple solution for continuous target replenishment, but suffer from mechanical instabilities that cause misalignment and reduced intensity. Cryogenic hydrogen targets (Gauthier et al., 2016; Obst-Huebl et al., 2017; Grieser et al., 2019; Polz et al., 2019) offer better spatial stability and higher flux, but are complex and costly. Micrometric solid targets, though widely used, still exhibit significant shot-to-shot variability, underscoring the need for improved alignment and target-holder design. Gas-jet targets (Jinno et al., 2013; Chen et al., 2017; Puyuelo-Valdes et al., 2019b) provide continuous operation without physical replacement and are easy to implement, but underperform in maximum energy and yield compared to solids, limiting their suitability for high-energy applications.

For many practical applications - including proton radiography (Romagnani et al., 2005), fast-ignition fusion research (Roth et al., 2001), warm dense matter and laboratory astrophysics (Patel et al., 2003; Albertazzi et al., 2014), neutron production (Roth et al., 2013; Higginson et al., 2015; Lelievre et al., 2024), medical therapy (Nemoto et al., 2001; Yogo et al., 2011), cultural heritage studies (Barberio et al., 2017b), injectors for larger accelerators (Antici et al., 2008, 2011), and materials science (Sciscio et al., 2017; Barberio et al., 2018a,b) - stability and repeatability outweigh peak performance. As a result, TNSA with solid foils remains the most reliable and reproducible solution.

However, because solid targets are ablated after every laser shot, high-repetition-rate TNSA inevitably relies on robust, automated multi-target positioning systems (Chagovets et al., 2021). Current setups can deliver sub-Hz multi-MeV beams (Gao et al., 2017; Treffert et al., 2022a; Giuffrida et al., 2023; Nedbailo et al., 2023), yet scaling beyond 1 Hz (Nishiuchi et al., 2009) remains challenging due to the dual requirements of fast target exchange and micron-level alignment. Overcoming these technical barriers is essential for transitioning next-generation laser-plasma accelerators from proof-of-concept experiments to practical, application-ready sources, motivating ongoing research on target delivery and shot-to-shot beam stabilization.

From a practical perspective, conventional accelerators are increasingly required to deliver precisely tailored combinations of beam current, energy, and quality (Rawlinson et al., 2002; Mathot et al., 2019). Laser-plasma accelerators offer a promising alternative, especially for material analysis, due to their flexibility and rapid switching capabilities. For example, they allow alternating between laser-driven PIXE and X-ray fluorescence (XRF) simply by changing the target's atomic number (Puyuelo-Valdes et al., 2021; Boivin et al., 2022). Beyond this versatility, laser-based ion

sources provide additional advantages over conventional systems : their broadband energy spectra allow real-time, depth-resolved analysis with energy selectors (Chen et al., 2014); their ultrafast temporal resolution (ns) enables observation of transient processes that conventional accelerators (ms) cannot resolve; and they support pump-probe spectroscopy via X-ray signals to track rapid electronic and atomic dynamics. However, these benefits are tempered by a generally lower ion flux, necessitating high-repetition-rate target refreshment to achieve meaningful yields (Passoni et al., 2019a; Mirani et al., 2021).

To overcome low ion yield troubles and shot-to-shot fluctuations, we developed an automated high-repetition-rate target system that ensures consistent target replenishment while preserving key advantages of laser-based ion sources, such as ultrafast resolution and versatile elemental analysis. We assess the repeatability and stability of TNSA-driven proton spectra at the 150 TW *ALLS* ion beamline using a multi-hundred (400) solid-target holder. Our analysis focuses on the stability of maximum proton energy, proton yield at various energies, energy-integrated proton yield, and laser-to-proton energy conversion efficiency. By systematically varying target materials and thicknesses, we evaluate the reproducibility of proton acceleration and benchmark our results against prior studies. This work is motivated by applications in PIXE and XRF (Passoni et al., 2019b; Barberio et al., 2019), which depend on stable proton beams in the 1–5 MeV range, critical for efficient X-ray production (Puyuelo-Valdes et al., 2021). Ensuring spectral stability in this window is essential for advancing laser-based spectroscopic techniques. Our high-repetition target system represents an important advance toward enhancing the reliability and practicality of laser-plasma ion sources, particularly for applications that demand stringent precision and consistent repeatability

3.2 Experimental setup

The experiment was conducted at the *ALLS* 150 TW ion acceleration beamline, located at *Institut national de la recherche scientifique* (*INRS*) in Varennes, near Montréal, Canada. The *ALLS* facility hosts a Ti : Sapphire 150 TW laser system (Vallières et al., 2020), operating at a maximum repetition rate of 2.5 Hz with a central wavelength of λ_0 = 800 nm. The laser system employs a double-Chirped Pulse Amplification (CPA) scheme, delivering laser pulses with an on-target energy of $\mathcal{E}_L$ = 3.2 J and an ultrashort pulse duration of τ_L = 22 fs. The experimental setup is illustrated in Figure 3.1.A. An f/3 off-axis parabola is used to focus the laser beam (measured at e^{-2}) down to a spot size of $w_{\mathrm{FWHM}} \sim 5\ \mu$m (Figure 3.1.B), achieving a peak intensity of $I_L \sim 1.3\times10^{20}$ W.cm^{-2}. To investigate intensity-dependent trends, we also conducted experiments at a lower peak intensity of $I_L \sim 3.3 \times 10^{19}$ W.cm^{-2}. To ensure a high-contrast laser pulse (Figure 3.2), a cross-wave polarizer (XPW) stage was implemented before the second CPA amplification. This setup effectively reduces Amplified Spontaneous Emission (ASE), achieving an ASE contrast <10^{-10} at -100 ps before the main pulse and a temporal contrast <10^{-8} at -10 ps and <10^{-6} at -3 ps, ensuring a steep intensity rise. The apparent prepulses observed around -20 ps before the main peak are not physical. These

features are known measurement artifacts intrinsic to the *Sequoia* third-order cross-correlator (Amplitude Laser, 2025a), typically arising from internal reflections and phase-matching imperfections in the nonlinear crystal. Therefore, they do not represent real pre-pulses and have no measurable influence on the laser–target interaction.

The p-polarized laser pulses were incident at 20° with respect to the target-normal, irradiating thin micrometric foils of various materials, including copper (Cu), tantalum (Ta), aluminum (Al), and polyethylene terephthalate (PET), at different thicknesses. These foils were mounted on a custom-designed multi-target holder (Figure 3.1.C) optimized for high-repetition-rate TNSA. Each 5×5 cm^2 target foil was secured between two stainless steel components forming the target holder, enabling up to 400 shots per pumping cycle at a maximum repetition rate of 0.625 Hz (a quarter of the laser maximum repetition rate). At repetition rates above 0.625 Hz, performance was limited by (i) drive–laser mis-synchronization and (ii) stage settling/position accuracy under continuous motion. Specifically, the target delivery system did not always reach the commanded position within the laser's firing window, so shots were occasionally taken off-target or during residual vibration. We assume that the dominant contributors were timing jitter between the laser's shot marker and the motion controller, variable settling times after high-acceleration moves, and intermittent encoder dropouts under vacuum/cable strain. Moreover, we operate in a multi-shot operation mode where all target positions were preloaded into the motorized stage control software (*OWIS* software shown in Figure 3.3), allowing the laser to be triggered synchronously at 0.625 Hz, fully compatible with the target holder's motion system. In this mode, alignment is verified only at the beginning of the shot sequence, as presented in Figure 3.4, typically on the target located in the bottom-left corner of the target holder, and is then kept fixed for the duration of the sequence. This strategy increases the likelihood of positional drift over time, which can progressively degrade the laser-target overlap and compromise beam stability. However, the use of a five-axis alignment system (three translational and two tilt degrees of freedom) helps mitigate this limitation by ensuring precise initial positioning. In contrast, single-shot mode offers a more precise alternative, where laser focus is readjusted before each laser shot on target, ensuring optimal conditions for every target. However, this method significantly reduces the shot repetition rate, limiting data acquisition efficiency.

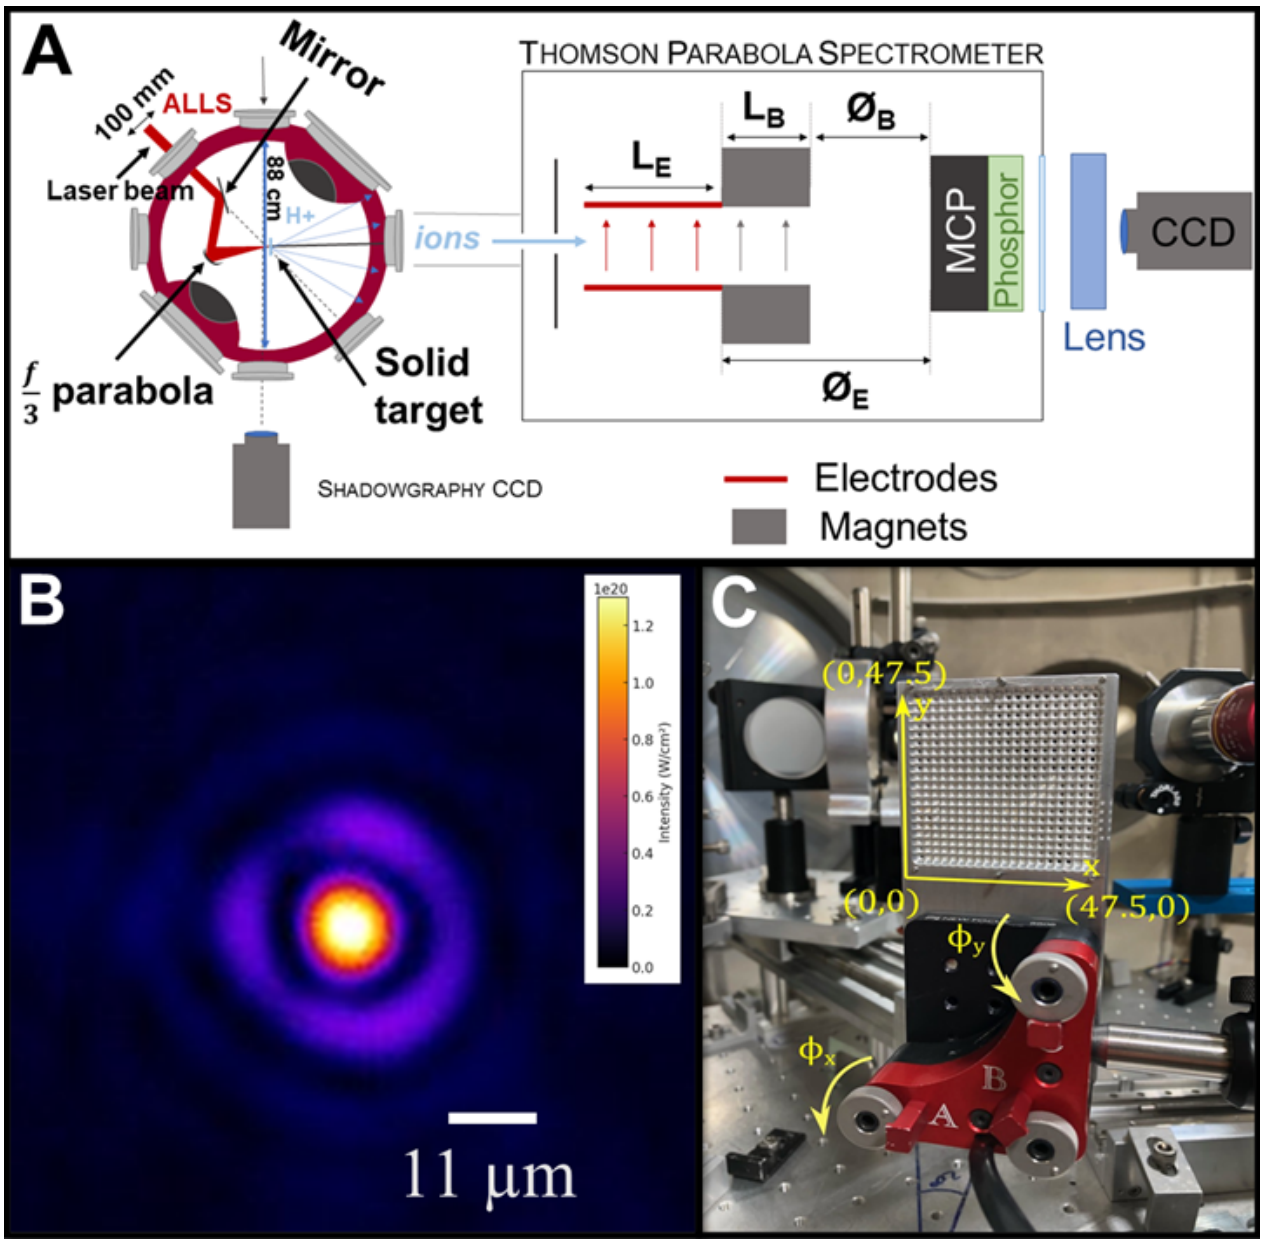

FIGURE 3.1 : Experimental setup, focal spot, and target holder

(A) Experimental chamber and Thomson Parabola spectrometer layout. (B) A view of the focal spot in the far field and (C) a picture of the custom-made target holder (dimensions in mm). This figure has been adapted from Catrix et al. (2023), with the publisher's permission.

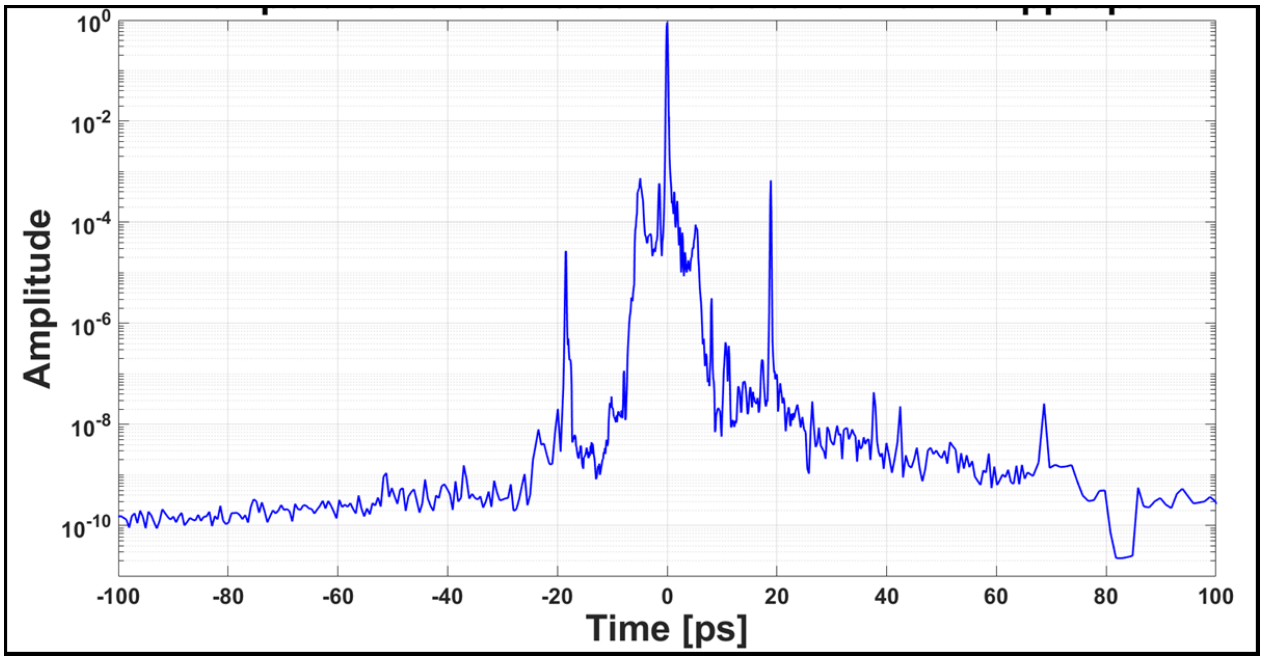

FIGURE 3.2 : Temporal contrast measurement of *ALLS* 150 TW laser.

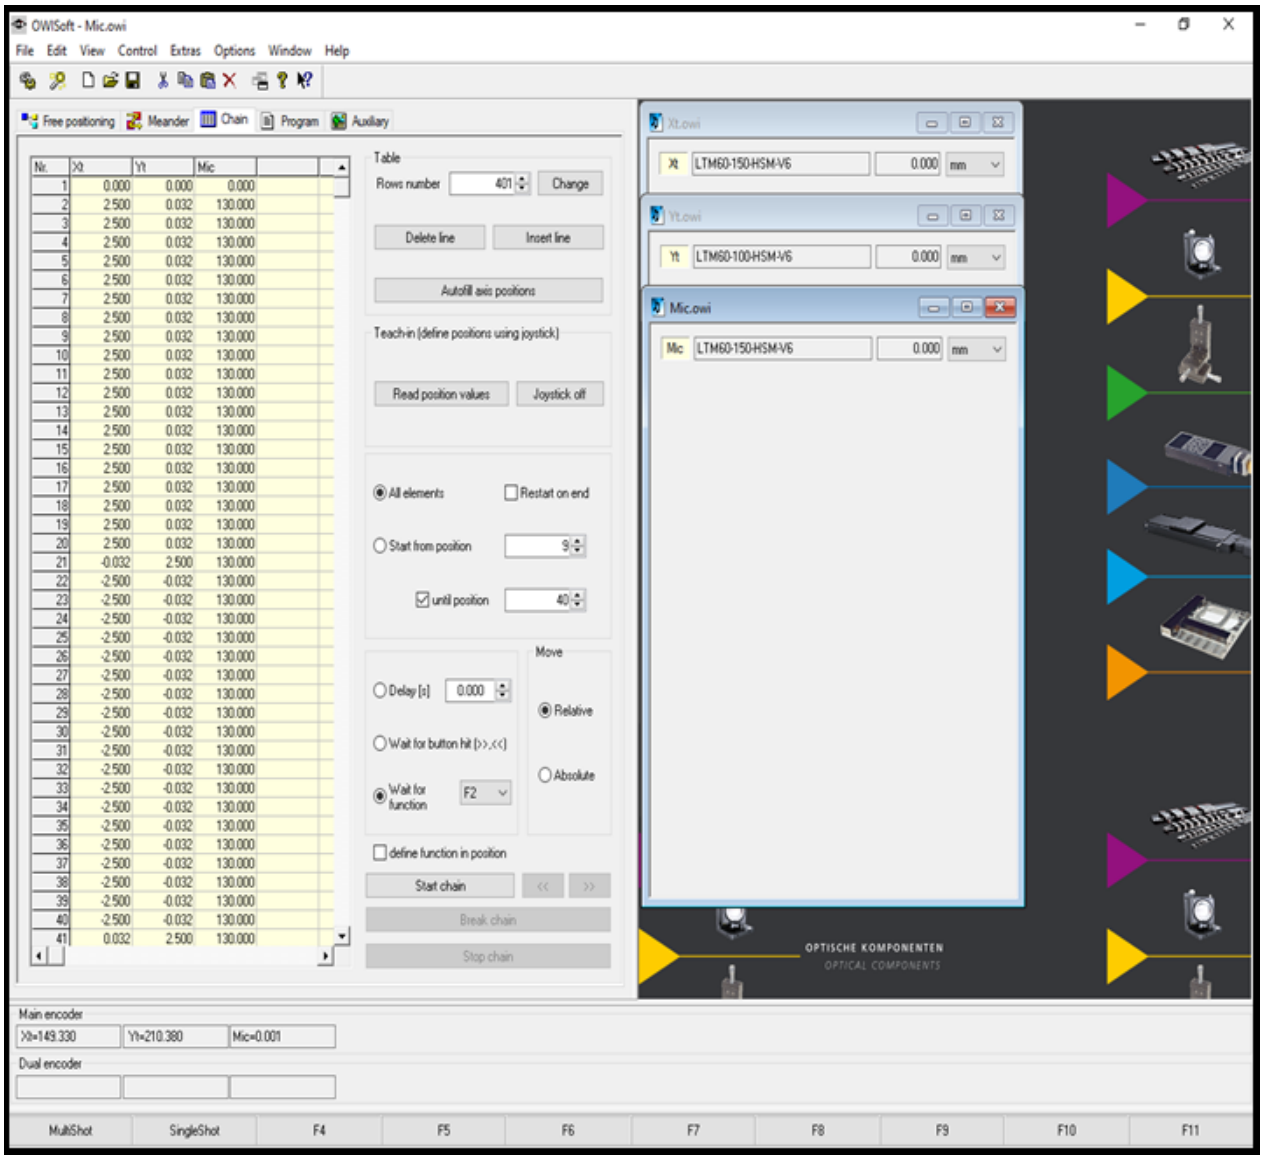

FIGURE 3.3 : *OWIS* chain-sequence control for multishot operation.

The chain function enables synchronized motion of the target-holder motors with the 0.625 Hz laser repetition rate, ensuring reliable shot-to-shot target refreshing during high-throughput experimental runs.

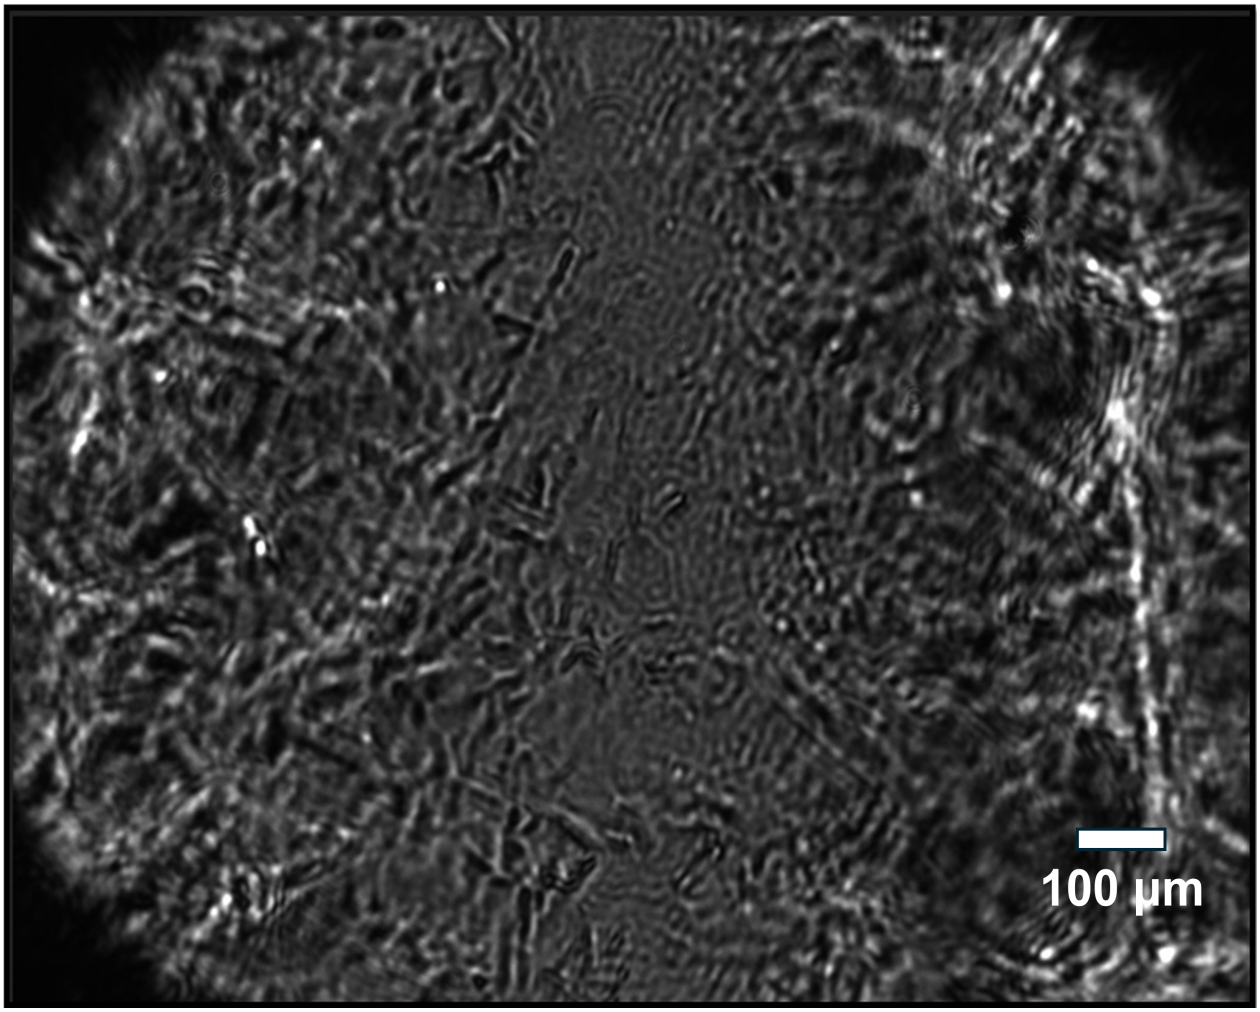

FIGURE 3.4 : Rear surface at best focus on a 2-μm-thick Ta foil.

The energy of the accelerated ions, including protons and carbon ions (ranging from C^{+} to C^{6+}), was characterized using a calibrated Thomson Parabola spectrometer (TPS) positioned along the target-normal axis (0°). The TPS was coupled with a MicroChannel Plate (MCP) detector, enabling the acquisition of ion trajectory images at every laser shot (Figure 3.5). A custom *MATLAB* script was developed to extract and plot the ion energy spectra from the raw data acquired by the TPS. This script plots and processes the ion spectra by averaging the signal intensity over multiple laser shots and quantifies shot-to-shot fluctuations by computing the standard deviation at each energy bin of ion beam metrics, providing a measure of the statistical dispersion across the ensemble of spectra. This study will focus solely on the analysis of proton beam characteristics, as protons are typically the dominant ion species accelerated via the TNSA mechanism due to their high charge-to-mass ratio and prevalence as surface contaminants (e.g., hydrocarbons or water vapor). Their lighter mass allows them to be accelerated to higher energies compared to heavier ions under the same sheath field conditions and therefore exhibit longer projected ranges, making them more relevant for in-depth analysis in applications such as radiobiology, material science, and medical physics. In this study, the maximum proton energy (i.e, the cut-off energy) $\mathcal{E}_{p,max}$ was determined by identifying the proton energy value located at the midpoint of the final slope in the spectrum. For further details regarding the experimental setup and laser system specifications, we refer the reader to the work : Vallières et al. (2020).

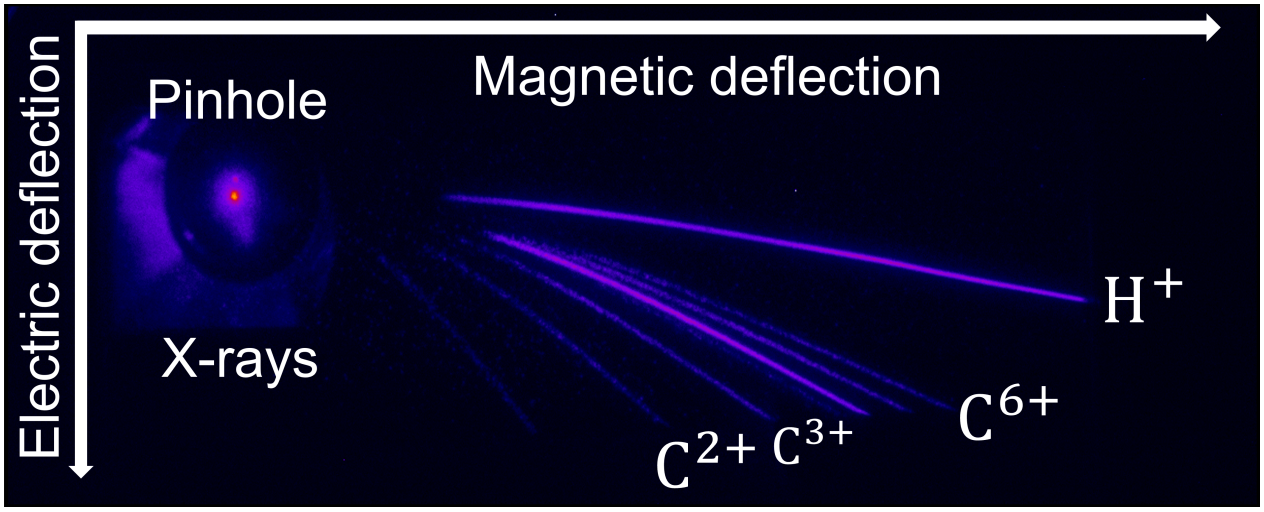

FIGURE 3.5 : Thomson Parabola image of ion traces.

A laser-triggered image from the Thomson Parabola spectrometer after irradiation of a 1-μm-thick copper foil. The parabola traces of the carbon species that were not considered in the analysis of the present study are also represented. In particular, the one on the left of C^{2+} is C^{+}, along with the two above C^{4+} being C^{5+} and C^{6+}.

To accurately analyze the images recorded on the MCP used within the TPS, it is essential to correctly track the ion trajectories, which vary according to their energy and charge-to-mass ratio under the influence of electric and magnetic fields in the TPS. Since these fields may not remain perfectly uniform along the ion paths, it is more appropriate to describe them using the field inte-

grals I_B (for the magnetic field) and I_E (for the electric field), introduced in equations 3.1 and 3.2. As highlighted by Ducret et al. (2018), this approach offers a more general characterization of the magnetic $\vec{B}$ and electric $\vec{E}$ fields, avoiding the oversimplified assumption of uniform, top-hat field profiles along the TPS axis. These integrals can then be used in the established non-relativistic, non-paraxial dispersion relations (equations 3.3 and 3.4), which describe the motion of charged particles through static fields in the TPS :

Equation 3.1 : Magnetic field integral in the TPS

$$I_B = \int_0^{L_B} \mathbf{B}(z)\,\mathrm{d}z = \overline{\mathbf{B}}\,L_B \tag{3.1}$$

Equation 3.2 : Electric field integral in the TPS

$$I_E = \int_0^{L_E} \mathbf{E}(z)\,\mathrm{d}z = \overline{\mathbf{E}}\,L_E \tag{3.2}$$

Equation 3.3 : Magnetic deflection of ions in the TPS

$$x(\mathcal{E}_K, I_B) = \frac{L_B}{q\,I_B}\left(\sqrt{2\mathcal{E}_K m - (qI_B)^2} - \sqrt{2\mathcal{E}_K m}\right) - \varnothing_B \frac{qI_B}{\sqrt{2\mathcal{E}_K m - (qI_B)^2}} \tag{3.3}$$

Equation 3.4 : Electric deflection of ions in the TPS

$$y(\mathcal{E}_K, I_E) = \frac{qI_E}{2\mathcal{E}_K}\left(\frac{L_E}{2} + \varnothing_E\right) \tag{3.4}$$

with the validity condition $2\mathcal{E}_K m \geq (qI_B)^2$. This condition corresponds physically to the requirement that the particle's kinetic energy is sufficiently large for it to traverse the magnetic field region without being completely turned back by the Lorentz force. In equations 3.3 and 3.4, the parameters q, m and $\mathcal{E}_K$ correspond to the particle's charge, mass, and kinetic energy, respectively. The quantities L_B and L_E denote the lengths of the magnetic and electric field regions, respectively, while $\varnothing_B$ and $\varnothing_E$ refer to the corresponding drift distances after each field region. A schematic representation is provided in Figure 3.1.A. These field-related distances are determined experimentally and are thus known (Table 3.1). Once the field integrals I_B and I_E are properly calibrated, the system becomes fully defined in terms of the particle's kinetic energy and its charge-to-mass ratio. With these calibrated parameters and the analytical expressions provided, the TPS can be used to quantitatively determine both the number and kinetic energy of laser-driven ions.

First, the magnetic field profile was measured using a calibrated gaussmeter. The field integral I_B was initially estimated using the magnetic field length L_B as a first approximation. For the electric field, the voltage difference ΔV between the electrodes and their separation d, as well as the field length L_E, was used to obtain a first estimate of I_E.

TABLE 3.1 : Electrostatic and magnetic parameters used in the TPS design.

Parameter	Description	Value
L_B	Magnetic field length	0.12 m
L_E	Electric field length	0.06 m
$\varnothing_B$	Magnetic field drift length	0.075 m
$\varnothing_E$	Electric field drift length	0.195 m
ΔV	Voltage difference between electrodes	15000 V
d	Electrode separation	0.02 m
B_y	Magnetic field strength	0.43 T
E_y	Electric field strength	7.5×10^5 V/m

To quantitatively analyze the traces obtained with the Thomson Parabola Spectrometer (TPS) coupled to a microchannel plate (MCP), we implemented a dedicated *MATLAB* routine that reconstructs the energy distribution of detected ions from their parabolic traces. The main steps of the analysis are as follows :

1. **Image preprocessing and smoothing :** A Gaussian convolution kernel of size $N_{\text{Gauss}} = 5$ and standard deviation $\sigma_{\text{Gauss}} = 1$ pixel is applied on the raw MCP image (Figure 3.5) to reduce high-frequency noise while preserving the parabolic trace structure :

   Equation 3.5 : Gaussian smoothing of the raw MCP image

   $$I_{\text{smooth}} = I_{\text{raw}} * G(x, y), \qquad G(x, y) = \exp\left[-\frac{x^2 + y^2}{2\sigma^2_{\text{Gauss}}}\right]. \tag{3.5}$$

   This filtering effectively removes shot noise without altering the spatial profile of the ion trajectories.
2. **Trajectory selection and trace extraction :** The code selects the ion species of interest (e.g., H^+, C^{2+}, C^{3+}, C^{4+}, ...) and applies a polynomial fit corresponding to the expected parabolic trajectory. The coefficients are initialized from previous experimental calibrations and refined iteratively to align with the observed trace.
3. **Signal projection and background subtraction :** The signal intensity is integrated along the parabolic trace over a finite transverse width of ± 5 pixels, corresponding to twice the 500 µm pinhole diameter. The background reference is taken from a triggered MCP of a shot in vacuum (without a target in the laser beam path), ensuring the subtraction of MCP dark current and scattered X-ray background.
4. **Energy calibration :** The transverse displacement of the ions recorded on the MCP is mapped to kinetic energy using equations 3.3 and 3.4.
5. **Spectrum reconstruction :** The differential number of particles per energy and solid angle $\frac{d^2N}{d\mathcal{E}d\Omega}$ (proton yield) is obtained from the intensity after applying a calibration factor and correcting for the detector's geometry and sensitivity. Energy bin width and angular divergence effects are included.

6. **Proton number and energy transfer calculation :** The total number of particles per steradian (energy-integrated yield) is computed via numerical integration of proton spectra. Additionally, the laser-to-proton energy conversion efficiency is estimated by integrating the total proton energy within the measured solid angle and normalizing by the laser energy.
7. **Validation :** The extracted experimental trace is compared to the theoretical parabolic trajectories derived from the known E and B fields, ensuring correct energy scaling. Overlaying multiple theoretical curves (H^+, C^{2+}, C^{4+}) confirms calibration consistency.

This approach provides an absolute, quantitatively calibrated characterization of laser-driven ion spectra, including accurate yield and conversion efficiency estimates for protons above minimum energy measured on the MCP.

3.3 Proton energy spectrum analysis and beam quality evaluation

Figures 3.6.A and B present the proton energy spectrum and the average proton spectrum, respectively, for a sequence of 296 consecutive laser shots irradiating a 4.5-μm-thick aluminum foil. These data serve as the basis for extracting key statistical parameters to assess the quality and stability of the produced proton beams. We selected this specific target thickness and material as it offers the best compromise between proton yield, maximum proton energy, and mounting stability. These factors are crucial when choosing targets for high-repetition-rate experiments, where maintaining consistent beam characteristics across multiple shots is essential. Figure 3.6.A reveals a gradual downshift in proton maximum energy over multiple shots, highlighting the effects of alignment drift that would not occur in single-shot mode. Despite this, the precision of the motorized stage and the optical imaging system minimizes energy fluctuations. Regardless of the alignment mode, this stability mainly stems from the $f/3$ off-axis parabolic focusing system, which ensures a pointing stability of 3.9 μm RMS laterally and 5.2 μm RMS vertically at the focal plane well within the 2-mm-diameter circular target size. From Figure 3.6.B, we inferred a hot electron temperature of approximately 0.9 MeV, estimated from the slope of the energy spectrum (Haines et al., 2009). The measured average maximum proton energy was (5.6 ± 0.4) MeV (i.e $\pm 7\%$), confirming the observed small gradual downshift in proton maximum energy in Figure 3.6.A. For the 4.5-μm-thick aluminum foil, the shot-to-shot fluctuation (standard deviation) over 296 laser shots in the measured proton yield was (Figures 3.6.A and .B) :

— 14.9% for the integrated proton yield,

— 18.2% for the proton yield at 1 MeV,

— 17.6% for the proton yield at 2 MeV,

— 25.7% for the proton yield at 4 MeV.

The observed shot-to-shot fluctuations in proton yield, integrated proton yield and maximum proton energy can be attributed to multiple factors :

— **Laser energy fluctuations** : Small variations in pulse energy directly affect the intensity at the focus (*ALLS* laser has a 2.5% RMS energy variation).

— **Long-term laser wavefront thermal drift** : Over time, thermal effects can introduce slow wavefront distortions, modifying the focal spot shape and energy distribution.

— **Target non-planarity** : Small deformations or surface irregularities in the target introduce inconsistencies in laser-plasma coupling from shot-to-shot, leading to fluctuations in the acceleration process (Brenner et al., 2011; Coury et al., 2012).

We then calculated the laser-to-proton energy conversion efficiency as :

Equation 3.6 : Laser-to-proton energy conversion efficiency

$$\alpha = \frac{\mathcal{E}_{\text{tot}}}{\mathcal{E}_L} = \frac{\int_{\mathcal{E}_{p,\min}}^{\mathcal{E}_{p,\max}} \frac{d^2N}{d\mathcal{E}\, d\Omega}\, \mathcal{E}\, d\mathcal{E}}{\mathcal{E}_L} \times \Omega. \tag{3.6}$$

where $\mathcal{E}_{tot}$ is the total proton energy, $\mathcal{E}_L$ is the energy of the laser, $\frac{d^2N}{d\mathcal{E}d\Omega}$ the proton yield, $\mathcal{E}_{p,min}$ the minimum proton energy detected by the MCP, $\mathcal{E}_{p,max}$ the maximum proton energy and Ω the solid angle covered by the conical proton beam. Our measurements indicate that $\alpha = 0.005$ defined by equation 4.5 remains relatively stable across multiple shots, with a shot-to-shot fluctuation of 12%.

To further explore how target properties influence laser-driven proton acceleration, we extended our study to include a range of target materials and thicknesses. We compared copper foils (1 μm and 3 μm), tantalum (2 μm), PET (1.5 μm), and the 4.5-μm-thick aluminum foil at a laser intensity of $I_L \sim 1.3 \times 10^{20}$ W.cm^{-2} and a maximum repetition rate of 0.625 Hz. As illustrated in Figure 3.6.C and reported in Table 3.2, comparisons across different target materials and thicknesses indicate that the maximum proton energy remained relatively stable across targets, with only modest variations. As expected, thinner targets yield slightly higher maximum proton energies, consistent with the scaling behavior typical of the TNSA regime (see Figure 1.A in Fuchs et al. (2006)).

TABLE 3.2 : $\mathcal{E}_{p,max}$ and α across foil types and thicknesses.

Foil type	Cu (Z=29) 1 µm	Ta (Z=73) 2 µm	Cu (Z=29) 3 µm	Al (Z=13) 4.5 µm
$\mathcal{E}_{p,\max}$ [MeV]	6.2 ($\pm$7%)	6.2 ($\pm$5%)	6.0 ($\pm$6%)	5.6 ($\pm$8%)
α [%]	0.7 ($\pm$21%)	1.0 ($\pm$16%)	0.6 ($\pm$15%)	0.5 ($\pm$12%)

A clear correlation emerges when comparing maximum energy with the laser-to-proton energy conversion and its shot-to-shot fluctuation : foils achieving higher energies tend to exhibit larger fluctuations, while thicker foils trade off peak energy for improved reproducibility. The most notable cases include :

— **2-μm-thick tantalum foil :** Achieved the highest maximum proton energy of approximately 6.2 MeV (Figure 3.6.C), with a low shot-to-shot fluctuation of 5.0% over 200 shots. This

stability is attributed to tantalum's intrinsic stiffness, which offers superior resistance to deformation compared to copper and aluminum.

— **1-μm-thick copper foil :** Produced a comparable maximum proton energy ($\sim$6.2 MeV), but with a higher standard deviation of 7.2% across 168 shots, indicating reduced shot-to-shot stability, likely due to the foil's lower mechanical robustness.

— **3-μm-thick copper foil :** Demonstrated improved beam stability with a standard deviation of 6.1% over 113 shots. The increased thickness likely enhances mechanical rigidity, mitigating alignment fluctuations observed with thinner copper foils. Reduced shot-to-shot stability, likely due to the foil's lower mechanical robustness compared to the tantalum foil.

While thinner foils generally support higher proton energies (Figure 3.7), our analysis reveals that they exhibit larger error bars considering the same material, primarily due to mechanical instability and increased alignment fluctuations. These issues affect the repeatability of the laser-target interaction, leading to higher variations in proton acceleration efficiency. Moreover, their effectiveness is contingent on maintaining an excellent pre-pulse contrast at this laser power level (150 TW) (Levy et al., 2008). Poor contrast can lead to premature plasma expansion, significantly altering the acceleration dynamics. Unlike metallic foils, PET targets introduce additional challenges due to their mechanical softness, which makes it difficult to maintain a perfectly flat surface at the laser focus. This misalignment leads to substantial fluctuations in proton energy, yield, and therefore beam quality. Additionally, they are transparent and weakly scattering light, making them hard to image or align with visible light diagnostics. In particular, our measurements for the 1.5-μm-thick PET foil yielded a maximum proton energy of 5.8 MeV, but with a very large shot-to-shot fluctuation of $\sim$45%, significantly higher than for metallic foils. These findings emphasize that, while PET foils remain an interesting alternative for specific applications, their use in high-repetition-rate laser-plasma experiments requires additional strategies to mitigate alignment instability and improve shot-to-shot consistency. Potential solutions could include pre-tensioning mechanisms or rigid support structures to reduce surface deformation.

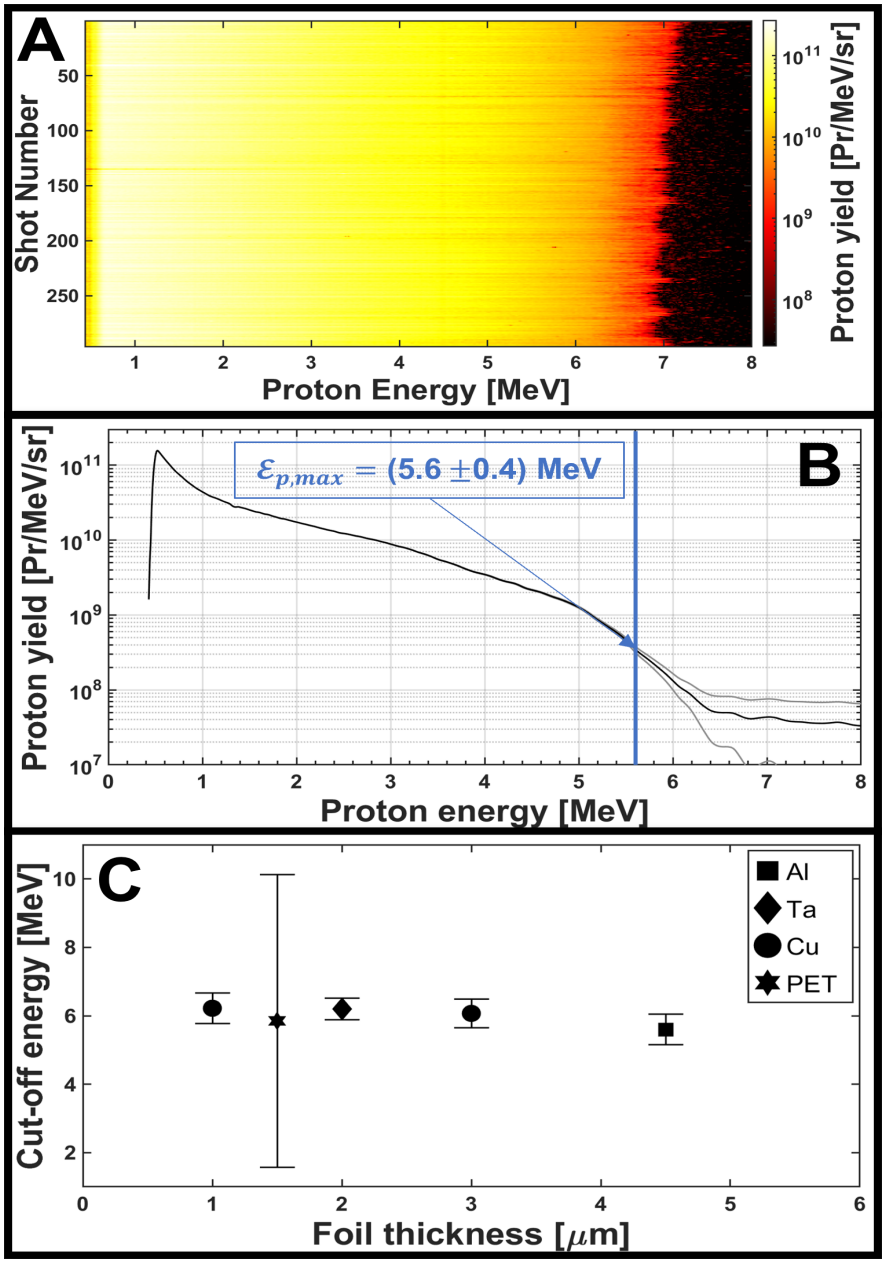

FIGURE 3.6 : Proton spectra and cut-off energy for various foils.

(A) Proton energy spectrum in the laser forward direction for 296 shots on a 4.5-μm-thick aluminum foil at a laser power of 150 TW recorded at 0.625 Hz. (B) Average spectrum of a sequence of 296 consecutive shots on a 4.5-μm-thick aluminum foil (the error bars are calculated using the square root of the minimum variance over the whole spectrum for better data visualization) (C) Mean maximum proton energy for a given foil irradiated by a laser intensity $I_L \sim 1.3 \times 10^{20}$ W.cm^{-2}; Shot-to-shot fluctuations are calculated using the standard deviation over 200 laser shots for tantalum 2 μm, 168 for copper 1 μm, 113 for copper 3 μm, 296 for aluminum 4.5 μm and 253 for PET 1.5 μm. These figures have been adapted from Catrix et al. (2023), with the publisher's permission.

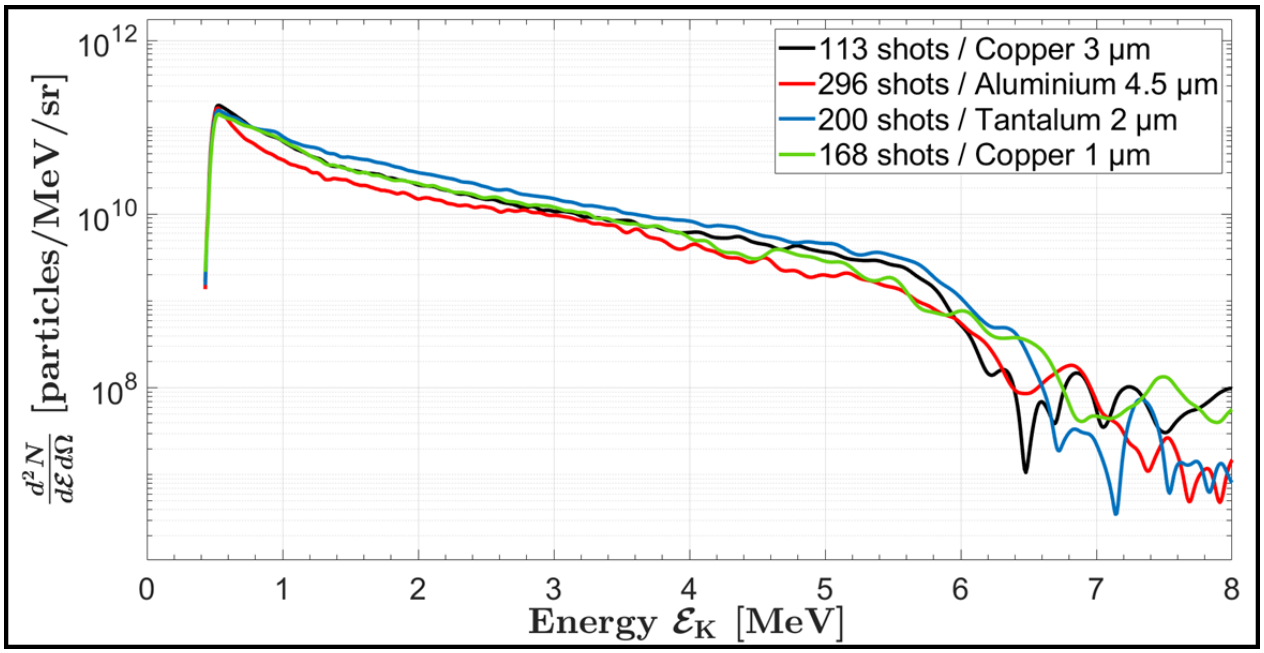

FIGURE 3.7 : Comparison of averaged proton spectra from different foil materials and thicknesses.

Shot-averaged proton spectra for different target configurations : 113 shots on a 3-μm-thick Cu foil, 296 shots on a 4.5-μm-thick Al foil, 200 shots on a 2-μm-thick Ta foil, 168 shots on a 1-μm-thick Cu foil. Error bars have not been included here for better visualization.

As shown in Table 3.3, the shot-to-shot fluctuation in maximum proton energy across all metallic foils in our experiment is 6.8%, indicating a high degree of stability relative to comparable studies conducted at other high-power laser facilities. A comparative study was conducted at the *ATLAS* 300 TW laser system (Gao et al., 2017), where researchers tested a multi-target holder using various target materials and thicknesses ranging from 50 nm to 1 μm — including 550-nm-thick polymer foils — with hexapod-controlled positioning. Their system operated at 0.5 Hz, slightly lower than our maximum 0.625 Hz repetition rate. However, their shot-to-shot variations were significantly larger, with a maximum proton energy fluctuation of $\sim$27% and a proton yield variation of $\sim$86% at 4 MeV. These results highlight the advantage of our setup in providing more consistent proton beam parameters over repeated shots. Further comparison with Green et al. (2014b) shows even larger fluctuations in maximum proton energy, with their study reporting variations greater than 15% for a 6-μm-thick aluminum foil and exceeding 30% for a 2-μm-thick aluminum foil. These results align with expectations, as thinner foils tend to exhibit larger fluctuations due to their sensitivity to pre-pulse contrast and laser energy variations. Another notable example is the *DRACO* laser facility in Germany, where high-repetition-rate proton beams were generated using a continuously flowing cryogenic hydrogen jet target (Gauthier et al., 2016). Their system achieved proton beams with an average maximum proton energy variation of $\sim$30%, while the proton yield at 3 MeV fluctuated by $\sim$50% from shot-to-shot. The use of cryogenic hydrogen jets provides an advantage in terms of target replenishment at high repetition rates, but it also introduces inherent fluctuations in density and positioning, contributing to the observed instabilities. Additionally, cryogenic hydrogen jets are often transparent or weakly scattering, making them hard to image or align with visible light diagnostics, while safety concerns arise from the flammability of hydrogen. Their implementation also requires sophisticated cooling systems (cryostats, vacuum insulation, cryo-compressors) to maintain hydrogen in liquid or solid form. Studies performed in Japan on the *J-KAREN* laser (Nishiuchi et al., 2009) have also reported similar shot-to-shot fluctuations, particularly for experiments involving quasi-monoenergetic lower-flux ion beams produced using passive magnetic elements. While their approach provides more controlled beam profiles, the trade-off is a higher degree of flux fluctuations compared to direct TNSA acceleration. At *ELI Beamlines* (Giuffrida et al., 2023), experiments have demonstrated the production of 10 MeV protons at a repetition rate of 0.5 Hz. However, their experimental setup is significantly more complex than ours, requiring additional beam transport and focusing optics, making direct comparisons challenging. Nonetheless, their results confirm the feasibility of high-energy proton production at high repetition rates, even though maintaining long-term stability remains a key challenge in such systems.

TABLE 3.3 : Averaged proton statistics at $I_L \sim 1.3 \times 10^{20}$ W.cm^{-2}.

Material (thickness)	$\mathcal{E}_{p,max}$ [MeV]	N_1 [Pr/MeV/sr]	N_2 [Pr/MeV/sr]	N_4 [Pr/MeV/sr]	N_{tot} [Pr/sr]	α [%]
Cu (Z=29) 1 µm Ta (Z=73) 2 µm Cu (Z=29) 3 µm Al (Z=13) 4.5 µm	6.0 (±6.8%)	6.1×10^{10} (±22.7%)	2.1×10^{10} (±18.8%)	5.0×10^{9} (±20.7%)	6.1×10^{11} (±18.7%)	0.5 (±12.4%)

Statistics at maximum laser intensity ($I_L \sim 1.3 \times 10^{20}$ W.cm^{-2}). The values are averages over 777 laser shots on metallic target foils and are extracted from the mean spectra. This table has been extracted from Catrix et al. (2023), with the publisher's permission.

* $\mathcal{E}_{p,max}$: Proton maximum energy.
* N_1 : Proton yield at 1 MeV.
* N_2 : Proton yield at 2 MeV.
* N_4 : Proton yield at 4 MeV.
* N_{tot} : Integrated proton yield.
* α : Laser-to-proton energy conversion efficiency.

For the integrated proton yield (N_{tot} [protons.sr^{-1}]), it appears that error bars on thinner targets are higher since they are mechanically less stable and subject to more alignment issues (Figure 3.8.B). For many practical applications, such as neutron generation (Lelievre et al., 2024) or PIXE spectroscopy (Puyuelo-Valdes et al., 2019a) via laser-driven ions, a specific energy range within the proton spectrum is more critical, often linked to the maximum cross-section of the nuclear or atomic process involved. We therefore quantify the proton yield and its stability at 1, 2 and 4 MeV (Figure 3.8.A) for the metallic foils. Although all points in Figure 3.8.A appear to overlap within their respective error bars; this does not necessarily imply that the target thickness does not influence the acceleration process. Rather, the relatively narrow range of foil thicknesses explored here (1–4.5 μm) and the observed shot-to-shot fluctuations, typically on the order of 15–25%, make it difficult to resolve a systematic trend. Nonetheless, its standard deviation for all types of foils does not overcome 25.7% (for the 4.5-μm-aluminum foil at 4 MeV). The 2-μm-thick tantalum foil produces the highest number of protons.sr^{-1} ($N_{tot} = 1.5 \times 10^{11}$ protons.sr^{-1} ± 18.6%) at $I_0 \sim 1.3 \times 10^{20}$ W.cm^{-2} because it is stiffer than any other material tested here, yielding a straighter foil on most laser shots, leading to a more efficient sheath acceleration. On *ATLAS* (Gao et al., 2017), a variation on the proton yield per laser shot of about 86% at 4 MeV was obtained. On the DRACO laser (Gauthier et al., 2016), the proton yield per laser shot at 3 MeV varies by roughly 50% from shot to shot. Relevant parameters covering all shots are found in Table 3.3 and exhibit better stability than previous work.

To further investigate the impact of laser intensity on proton acceleration, we conducted experiments at a reduced intensity of $I_L \sim 3.3 \times 10^{19}$ W.cm^{-2}, enabling us to assess energy scaling and stability under weaker field conditions. As shown in Figure 3.9.A, the mean maximum proton

energy remains relatively uniform across different materials and thicknesses, ranging from 1.72 MeV (1 μm Cu) to 1.90 MeV (2 μm Ta and Al), in agreement with the TNSA scaling law exhibited in Borghesi et al. (2008). The standard deviation in maximum proton energy across target types is in the same range as in high-intensity regimes ($I_L \sim 1.3\times10^{20}$ W.cm^{-2}), ranging between 5% and 9%. However, Figure 3.8.B reveals that N_{tot} exhibits greater shot-to-shot variability at $I_L \sim 3.3 \times 10^{19}$ W.cm^{-2}. This latter can be attributed to :

— **Enhanced sensitivity to surface imperfections** : The high temporal contrast of the *ALLS* laser suppresses preplasma formation, creating a sharp density gradient at the critical surface. Without a preplasma layer, due to the lower laser intensity, microscopic surface irregularities are not smoothed out, and the incident laser couples directly to the solid-density target. Consequently, local variations in morphology (e.g., roughness, contaminants, or defects) can significantly affect absorption, modify the sheath-field distribution, and thus influence the acceleration efficiency. Similar behavior has been discussed in a study on laser-induced breakdown and surface conditioning, which emphasizes the role of incipient plasma formation in homogenizing absorption and reducing shot-to-shot variability (Yang et al., 2022). Future investigations could systematically explore this regime by combining controlled preplasma generation (e.g., using a low-intensity prepulse or external preionization) with time-resolved diagnostics to directly correlate surface morphology, preplasma scale length, and the resulting proton-beam stability.

Since the maximum proton energies remain below 2 MeV, we evaluated proton yield at 1 MeV (Figure 3.9.B). Its shot-to-shot fluctuation mirrors that of the total yield, confirming that :

— The majority of protons are concentrated around 1 MeV, making it a relevant stability evaluation metric.

— Material- and thickness-induced yield fluctuations are more significant at lower intensities.

3.3.1 Discussion

Overall, mechanical instabilities in plastic (PET) targets, particularly under multi-shot operation, compromise reproducibility and alignment, making them unsuitable for precision applications. For scenarios requiring both high yield ($> 10^{11}$ protons.sr^{-1}) and energy, high-Z target materials (Ta or Cu) are generally more effective for generating high-energy protons in TNSA because of their enhanced laser absorption mechanisms, such as Brunel absorption or resonance absorption in pre-plasma. This results in greater energy transfer from the laser to hot electrons. Moreover, the electron stopping range in high-Z materials is shorter, which causes hot electrons to accumulate near the rear surface. This builds a denser and stronger sheath, leading to more efficient ion acceleration and higher maximum proton energies. In contrast, the 4.5-μm-thick aluminum foil, while more stable regarding the integrated proton yield, produces fewer hot electrons and therefore generates a weaker sheath field, demanding significantly more shots to reach equivalent yields :

— 30% more than 1 μm Cu,

— 70% more than 2 μm Ta,

— 40% more than 3 μm Cu.

Despite this, aluminum's superior mechanical stability supports long acquisition sequences with minimal yield and energy drift. The setup also achieves a stable laser-to-proton conversion energy efficiency of (0.5 ± 0.1)%.

Therefore, this study demonstrates that careful target optimization enables reliable, stable proton beams, making the approach well-suited for high-repetition-rate applications demanding efficiency and reproducibility.

Future advances toward higher repetition rates could be achieved through tighter synchronization and active feedback control :

— **Hardware synchronization :** Phase-lock the motion controller to the laser's reference clock and distribute triggers through low-jitter hardware lines.

— **Two-step handshake :** Implement a logical interlock in which the motion controller asserts a `POSITION_READY` signal only when the target is stationary within tolerance (e.g., ± 10-20 μm) before the laser gate is opened.

— **Predictive motion scheduling :** Estimate settling times dynamically and initiate target motion early so that the stage is fully stabilized before each shot.

— **Live position monitoring :** Integrate a fast imaging or displacement sensor to verify the target alignment in real time, allowing the system to delay or inhibit a shot if misalignment is detected.

With these improvements, particularly hardware-level synchronization and real-time feedback, the system could realistically approach full exploitation of *ALLS* laser maximum repetition rate of 2.5 Hz while maintaining high shot-to-shot reproducibility.

3.4 Application to PIXE spectroscopy

As previously mentioned, the 4.5-μm-thick aluminum foil offers an optimal trade-off between proton yield, maximum energy, and mechanical stability, making it a reliable choice for laser-driven PIXE. The effectiveness of PIXE relies on proton energies within the 1–5 MeV range, which are ideal for inducing characteristic X-ray emission from many elements, especially when proton/photon-to-X-ray conversion efficiency is high (e.g., metals). When combined with laser-driven XRF, the technique, referred to as XPIF (X-ray and Particle-Induced Fluorescence), provides complementary information for comprehensive material analysis, particularly in trace element identification. However, aluminum's K-shell emission energies ($K_\alpha \approx 1.49$ keV, $K_\beta \approx 1.56$ keV) fall below the lower detection limit of our X-ray detector, which is typically sensitive in the 2–25 keV range. As such, aluminum is not suitable for generating detectable characteristic X-rays in laser-driven XRF,

limiting its use in full XPIF characterization. A notable advantage of laser-driven ion accelerators is their versatility : by simply replacing the target with a high-Z material such as tantalum, the same setup can be readily adapted from PIXE to full XPIF operation. Tantalum's L-shell emission energies ($L_\alpha \approx 8.1$ keV, $L_\beta \approx 9.3$ keV) lie well in the detection range. These emissions enable XPIF measurements, making tantalum an effective choice for extending the analytical capabilities of the setup without substantial modifications to the laser or beamline parameters.

In many cases, a single laser shot can generate sufficient proton and photon flux to produce a clear X-ray spectrum, allowing for the identification of material components in a sample (Barberio et al., 2017b). However, certain materials, particularly organic materials, wood, and composite structures, exhibit lower proton/photon-to-X-ray conversion efficiency, making it difficult to achieve an X-ray peak that meets identification specifications. In such cases, accumulating multiple laser shots is the only viable method to enhance the signal-to-noise ratio (SNR) and improve detection sensitivity. To illustrate the importance of stable and cumulative high-repetition-rate proton beams in XPIF, Figure 3.10 presents a comparison of X-ray spectra obtained from a Madder Lake Mock-up, a wood finish historically used in Cremonese (Italy) bowed string instruments. This sample was analyzed using XPIF with a 2-μm-thick tantalum target, which supports both PIXE and laser-driven XRF due to its higher proton maximum energy and efficient X-ray production. Here are the key observations from the X-ray emission spectra (Figure 3.10) :

— **Single laser shot :** The arsenic (As) peak at 10.54 keV is not detectable, indicating that a single-shot signal is insufficient for clear peak identification.

— **Accumulation over 50 laser shots :** The arsenic peak becomes visible, demonstrating that multi-shot accumulation significantly improves the SNR, making it possible to detect trace elements in historical materials.

This example highlights the importance of stable, high-repetition-rate proton beam production for analytical techniques demanding high signal quality. By optimizing target material and ensuring consistent proton flux, XPIF becomes a powerful tool for historical material analysis, cultural heritage preservation, and advanced compositional diagnostics. These results emphasize the need to tailor target selection and refine multi-shot strategies to enhance X-ray signal detection for materials requiring enhanced X-ray signal detection.

3.5 Conclusion

Using a custom-designed multi-target holder, we have demonstrated that the *ALLS* 150 TW ion acceleration beamline can reliably generate stable proton beams (and therefore ion beams) over hundreds of consecutive laser shots in the TNSA regime. Despite inherent shot-to-shot fluctuations, the proton yield and maximum energy remain highly consistent, making this setup well-suited for high-repetition-rate applications such as laser-based PIXE and XPIF. Thinner metallic foils — such as 2-μm tantalum and 1-μm copper — produce higher-energy protons and greater yields,

whereas 4.5-μm aluminum offers the best long-term operational stability, which reduces alignment errors and ensures better overlap with the laser over hundreds of shots. In contrast, PET foils exhibit significant fluctuations, limiting plastic target use in precision contexts. The robustness of this configuration, combined with fast alignment, enables efficient multi-shot sequences with minimal drift, enhancing the reliability and repeatability of laser-based ion sources. These results represent a significant improvement over previous stability benchmarks and underscore the potential of this approach for advancing compact, high-repetition-rate laser–plasma accelerators. Ensuring such stability is critical for further beamline optimization and for enabling other high-precision applications.

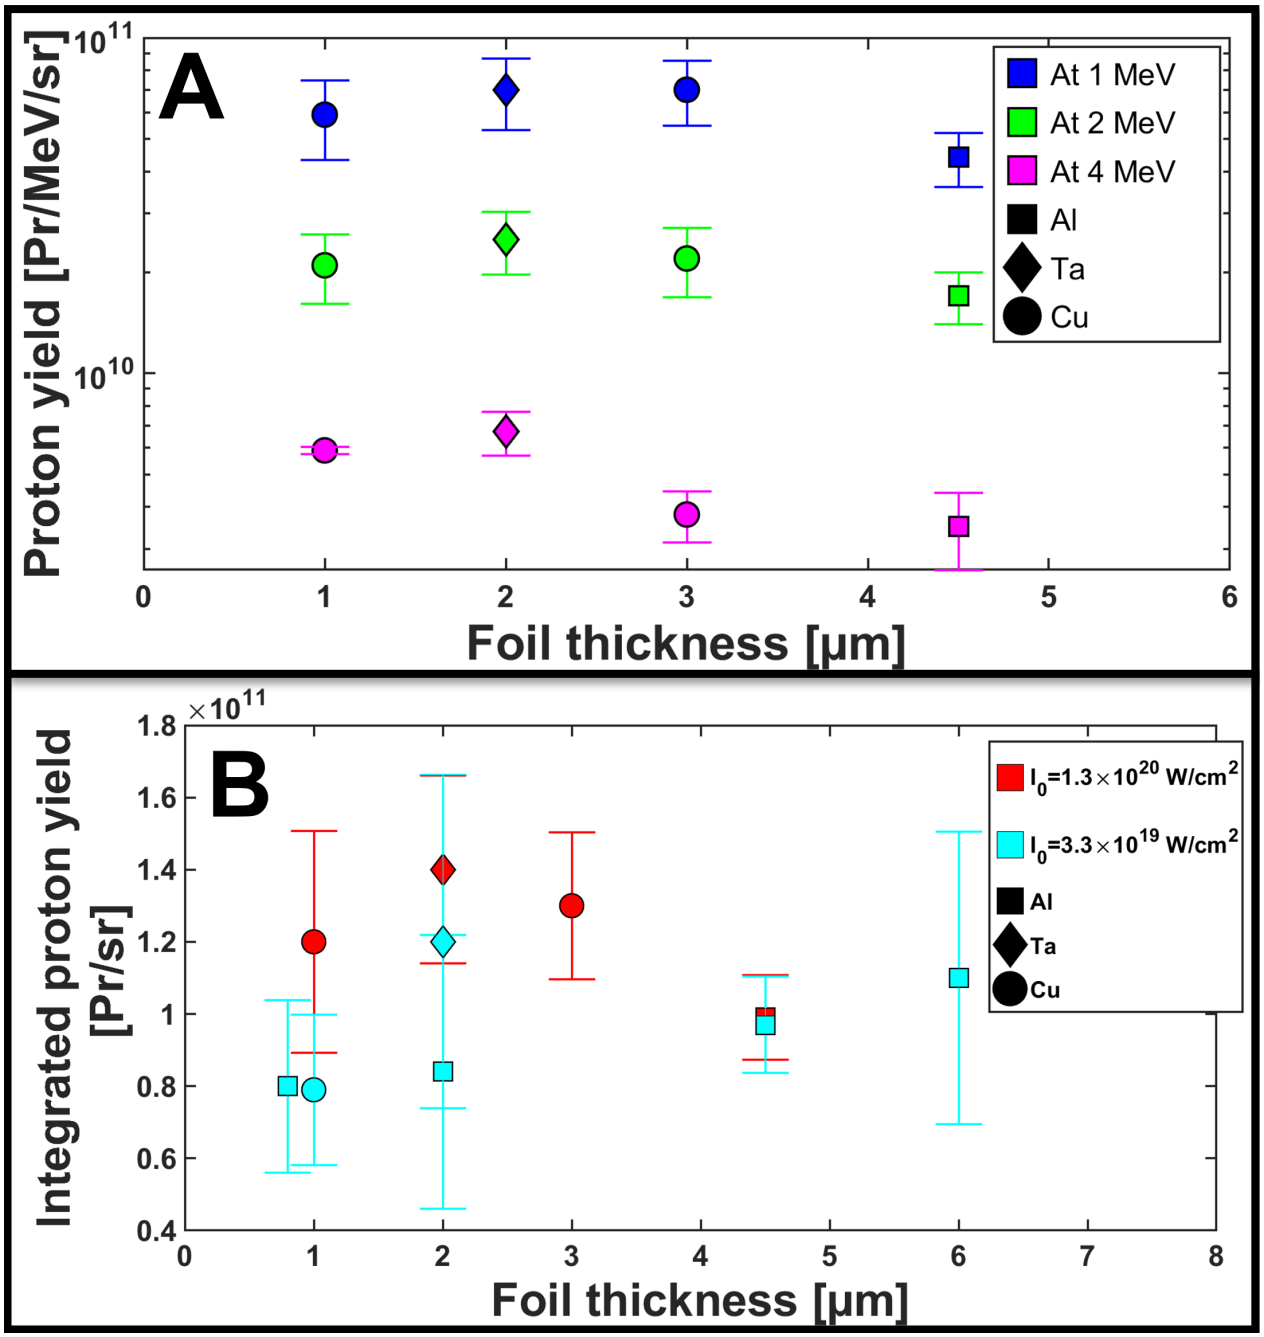

FIGURE 3.8 : Proton yield versus foil thickness at two intensities.

(A) Mean proton yield at 1, 2, and 4 MeV (protons/MeV/sr) at laser intensity $I_L \sim 1.3 \times 10^{20}$ W.cm^{-2}. (B) Total number of protons per steradian [protons.sr^{-1}]. Variations are calculated using the standard deviation over 200 laser shots for tantalum 2 μm, 168 for copper 1 μm, 113 for copper 3 μm, and 296 for aluminum 4.5 μm. These figures have been adapted from Catrix et al. (2023), with the publisher's permission.

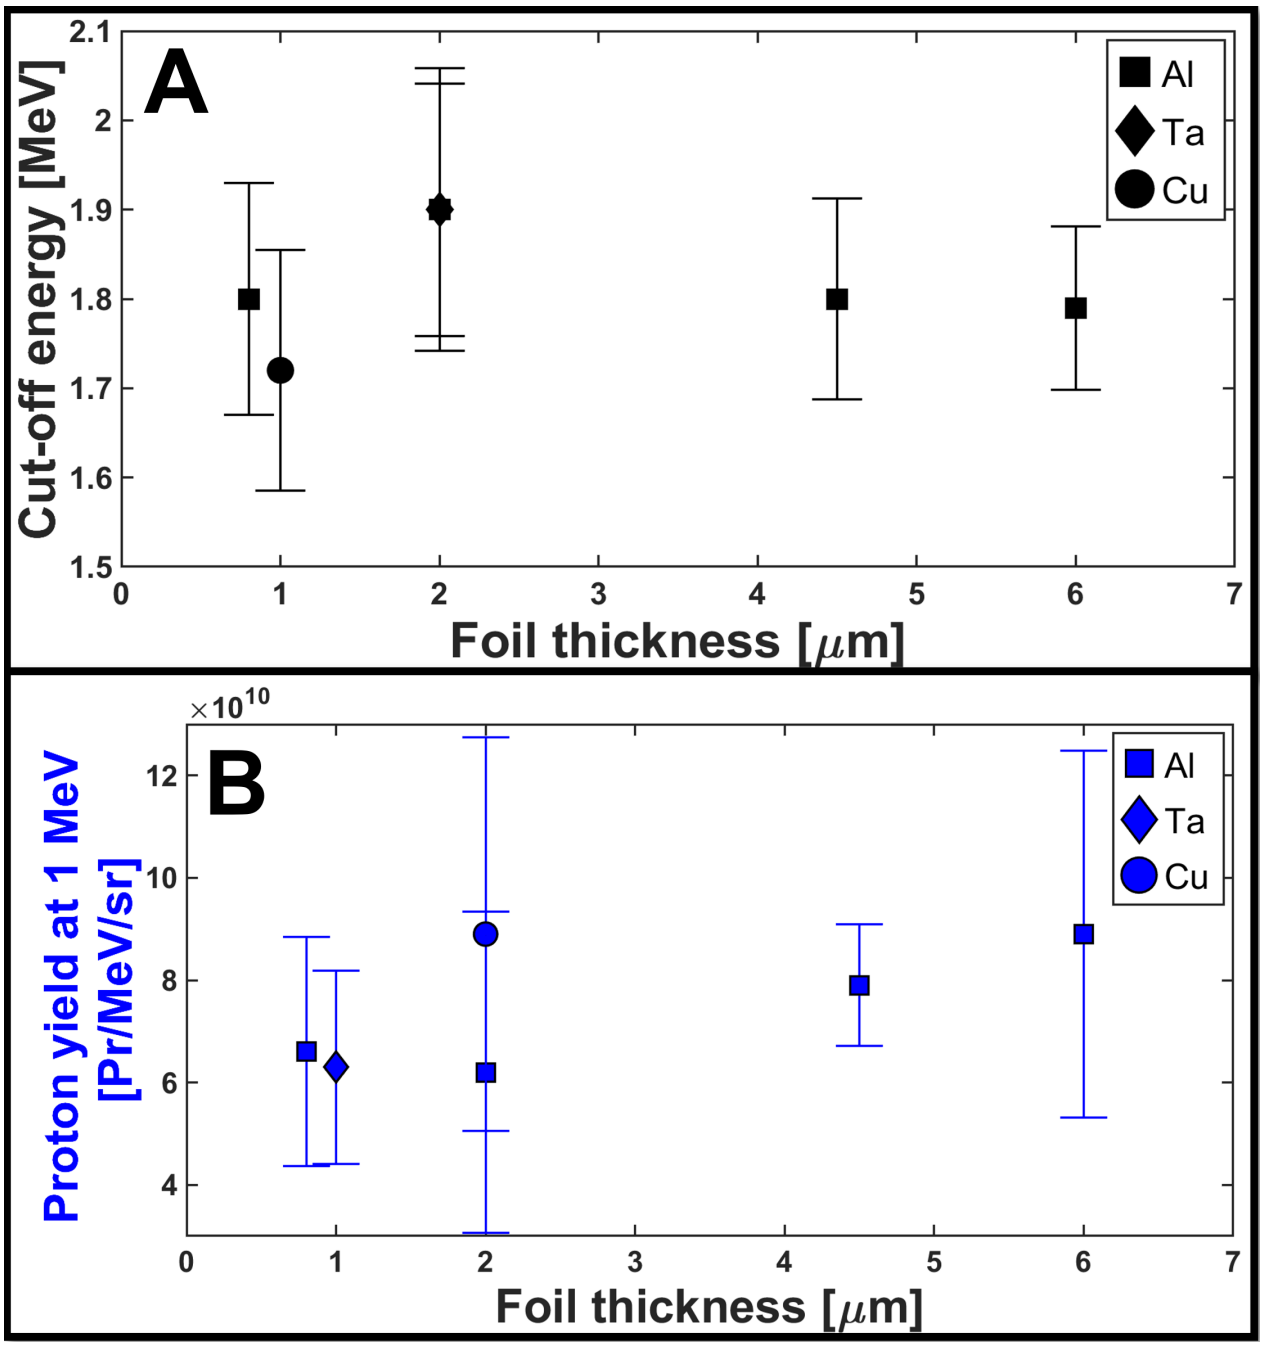

FIGURE 3.9 : Cut-off energy and proton yield at $I_L \sim 3.3 \times 10^{19}$ W.cm^{-2}.

(A) Mean maximum proton energy at laser intensity $I_L \sim 3.3 \times 10^{19}$ W.cm^{-2}; (B) Proton yield at 1 MeV at laser intensity $I_L \sim 3.3 \times 10^{19}$ W.cm^{-2}. Variations are calculated using the standard deviation over 200 laser shots for tantalum 2 μm, 168 for copper 1 μm, 113 for copper 3 μm and 296 for aluminum 4.5 μm. These figures have been adapted from Catrix et al. (2023), with the publisher's permission.

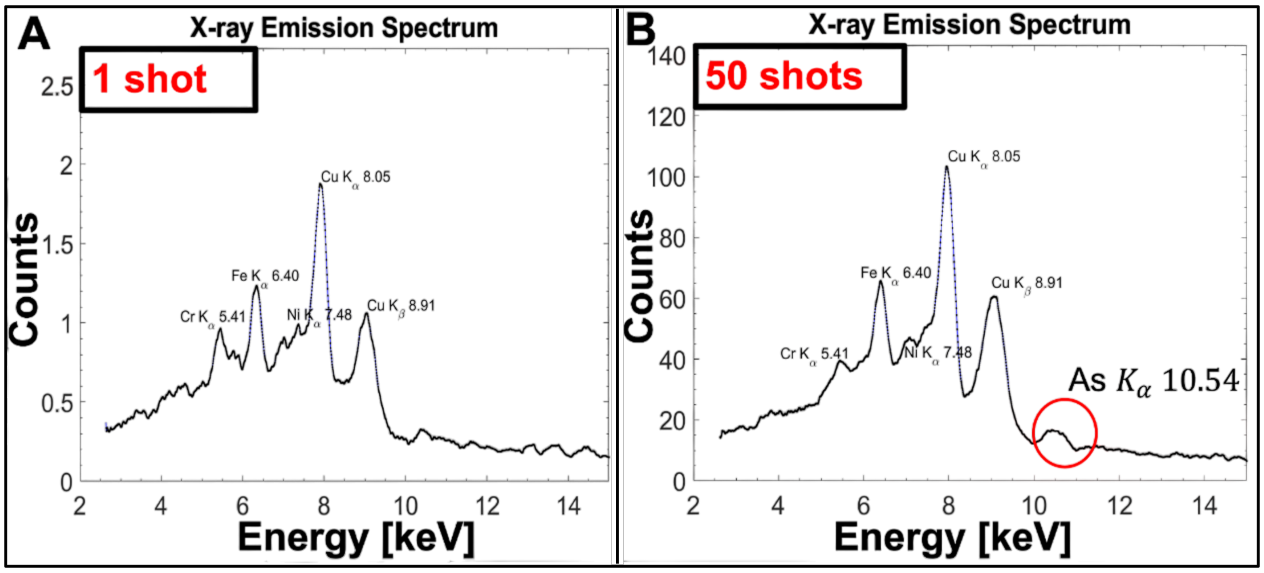

FIGURE 3.10 : X-ray spectra of a Madder Lake Mock-up.

(A) X-ray emission spectrum of a Madder Lake Mock-up (piece of wood) for one laser shot; (B) X-ray emission spectrum of a Madder Lake Mock-up (piece of wood) cumulated over 50 consecutive laser shots. These figures have been adapted from Catrix et al. (2023), with the publisher's permission.

4 PARAMETRIC CONTROL OF LASER-DRIVEN PROTON BEAMS THROUGH LASER SPATIO-TEMPORAL SHAPING

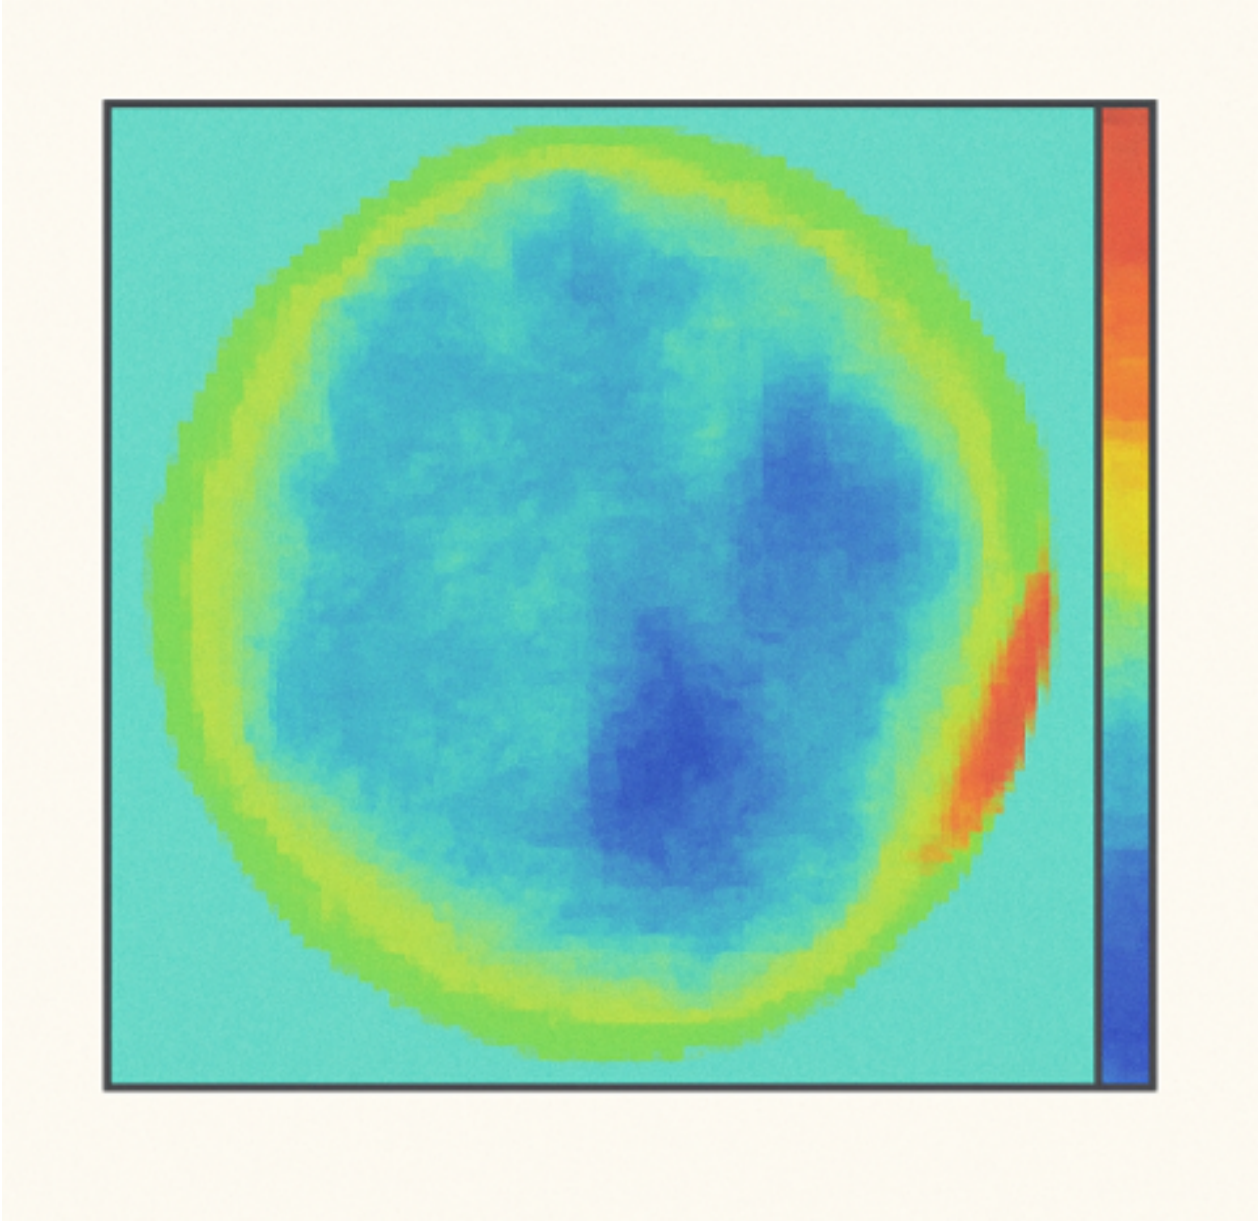

Spatial phase distribution of the laser beam. Original image created by the author (Elias Catrix). No external permission is required; use authorized by the author.

Foreword — This chapter presents a study of laser-driven proton acceleration using spatio-temporal laser beam shaping in the Target Normal Sheath Acceleration (TNSA) regime. By employing 20 of the 48 actuators of a deformable mirror in combination with an acousto-optic programmable dispersive filter, we independently tailored the laser's spatial wavefront and spectral phase. Using 4.5-μm-thick aluminum targets, spatial shaping had a pronounced influence on proton beam quality, whereas temporal shaping produced only a marginal effect. We further identify the maximum proton energy as a strong predictor of overall TNSA performance, establishing it as an effective single-parameter lever for the global optimization of laser-driven proton sources.

4.1 Introduction

Stable and efficient particle acceleration in high-intensity laser–plasma interactions relies on precise control of the laser's spatial and temporal properties, which govern hot-electron generation and transport. The properties of these hot electrons, especially their temperature and the steepness of local density gradients, ultimately control sheath strength and the resulting ion acceleration (Wilks et al., 2001). The spatial properties of the laser, such as width, divergence, propagation, and wavefront quality, are critical for achieving optimal focusing and efficient energy deposition (Fourmaux et al., 2008). Additionally, temporal characteristics, including pulse duration, spectral phase, and temporal contrast, also influence laser–matter coupling (Kaluza et al., 2004; Fuchs et al., 2006; Fourmaux et al., 2011b; Dollar et al., 2013). This temporal structure dictates the likelihood of pre-plasma formation, often triggered by pre-pulses or pedestal intensity, which can affect hot electron energy absorption efficiency by modifying front-surface absorption and weakening the rear-surface sheath field. Complementarily, target characteristics, such as composition, thickness, surface cleanliness, and geometry, influence ion beam divergence, maximum energy, and overall quality (Vallières et al., 2018; Catrix et al., 2023). Capturing these spatio-temporal dependencies requires diagnostics capable of resolving both the temporal amplitude and phase of ultrashort laser pulses (Trebino et al., 1997; Iaconis et al., 1998; Amplitude Laser, 2025b) as well as the laser wavefront (Phasics, 2025). Such comprehensive measurements are essential for understanding and controlling the mechanisms driving TNSA, ultimately enabling precise ion-acceleration dynamics and consistent high-quality beam generation.

Building on these diagnostic insights, adaptive optics - particularly deformable mirrors (DMs) - are commonly used to correct wavefront distortions and optimize the laser focus (Lei et al., 2012; Lefaudeux et al., 2013; Holmes, 2022). By adjusting actuator voltages, DMs enable real-time correction of aberrations and can also be used to tailor the wavefront for specific interaction conditions deliberately. When combined with spectral phase control using, for example, an acousto-optic programmable dispersive filter (AOPDF) in high-power laser experiments, this integrated approach allows systematic tuning of the laser in both space and time.

Together, these control tools set the stage for Machine learning (ML) approaches to efficiently explore and optimize TNSA performance. Neural network–based surrogate models have predicted maximum proton energy and ion yield from laser intensity, focal position, target thickness, and pre-plasma scale length, enabling rapid exploration of high-dimensional parameter spaces (Djordjević et al., 2021b; Schmitz et al., 2023). Stacked convolutional neural networks–Recurrent neural networks models extended this to ion acceleration time, emphasizing the role of temporal pulse shaping and pre-plasma gradients (Djordjević et al., 2021a), while neural networks with uncertainty quantification predicted hot-electron generation from pulse-shaping parameters (Mariscal et al., 2024). Moreover, statistical and Bayesian methods have linked laser energy, pulse duration, and hot-electron properties to maximum proton energy and efficiency, with multivariate regression providing predictive scaling laws above 100 MeV (Takagi et al., 2021; Desai et al., 2024). Moving to online optimization, a closed-loop Bayesian optimization (BO) simultaneously tuned five Zernike-mode wavefront aberrations and the target position relative to best focus to maximize proton beam energy in real-time, achieving performance comparable to manual tuning with only 57% of the laser energy (Loughran et al., 2023). Collectively, ML methods connect controllable laser and target parameters to key TNSA performance metrics.

In this chapter, we systematically investigate the influence of spatio-temporal beam shaping on laser-driven proton acceleration in the TNSA regime. By decoupling spatial and temporal degrees of control, we evaluate their individual impact on key proton beam properties. Spatial shaping is performed by adjusting 20 out of 48 actuators of a bimorph DM, while temporal shaping involves tuning the Group Delay Dispersion (GDD, β_2 [fs^2]) and the Third-Order Dispersion (TOD, β_3 [fs^3]) (Ghigo et al., 2007) spectral phase terms using an AOPDF. Using a combination of experimental measurements and numerical modeling with 4.5-μm-thick aluminum targets, we quantify how each component influences laser-driven proton beam metrics. Importantly, our analysis reveals that under our experimental conditions, maximum proton energy is the most sensitive and informative metric that exerts the greatest influence on acceleration efficiency. This study thus clarifies the distinct roles of spatial and temporal shaping and supports the use of maximum energy as a reliable single-parameter metric for guiding optimization strategies.

4.2 Deformable mirror actuators for wavefront control

DMs are a fundamental technology in adaptive optics, usually designed to correct wavefront distortions in high-precision optical systems. They consist of a thin reflective membrane whose surface can be actively controlled by an array of actuators. These actuators induce localized deformations, enabling real-time control of the mirror surface essential for precision applications such as astronomical imaging, microscopy, and high-power laser systems. DM actuators can be electrostatic (Bifano, 2011), piezoelectric (Kokorowski, 1979; Toporovskiy et al., 2019a), or magnetic (Biasi et al., 2010), each offering distinct advantages in terms of response time, stroke range, and preci-

sion. The piezoelectric DMs in turn could be divided into two large categories : bimorph (Zhu et al., 2017; Kazasidis et al., 2018; Toporovskiy et al., 2019b) as the one available on the *Advanced Laser Light Source* (*ALLS*) 150 TW ion beamline and stacked-actuator DMs (SADMs) (Wlodarczyk et al., 2014; Samarkin et al., 2021).

4.2.1 Structure and working principle of deformable mirrors

A DM is composed of several key elements :

— **Reflective membrane** : A thin, highly polished metal-coated surface that reflects incoming light. This membrane must be both flexible and resilient to withstand frequent shape changes.
— **Actuator array** : Distributed beneath the membrane, these actuators push or pull the surface to compensate for wavefront distortions.
— **Control system** : A feedback mechanism that computes the necessary corrections based on wavefront sensor (WFS) data and adjusts the actuators' pushing forces accordingly.
— **Support structure** : A base that holds the actuators and provides mechanical stability.

A sketch of a typical bimorph DM is shown in Figure 4.1. The deformation of the mirror surface allows correction of wavefront errors (see Figure 4.2), ensuring an optimal focal spot and reducing distortions due to atmospheric turbulence or optical aberrations in telescopes for example.

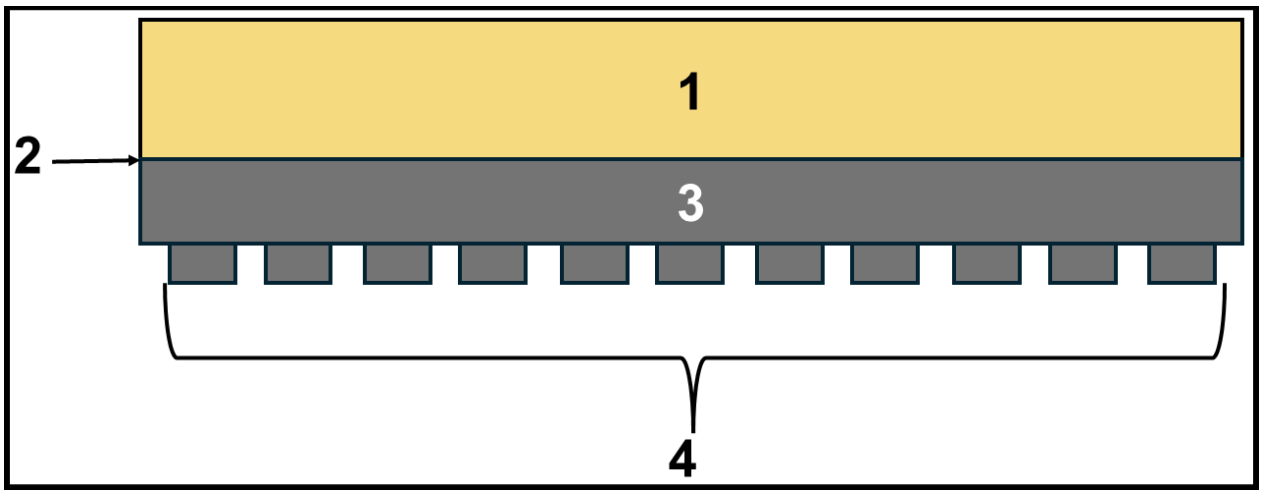

FIGURE 4.1 : Structure of a bimorph deformable mirror.

1 : Substrate with a reflective coating = Reflective membrane, 2 : Grounding electrode, 3 : Piezoceramic disk, and 4 : Actuator array.

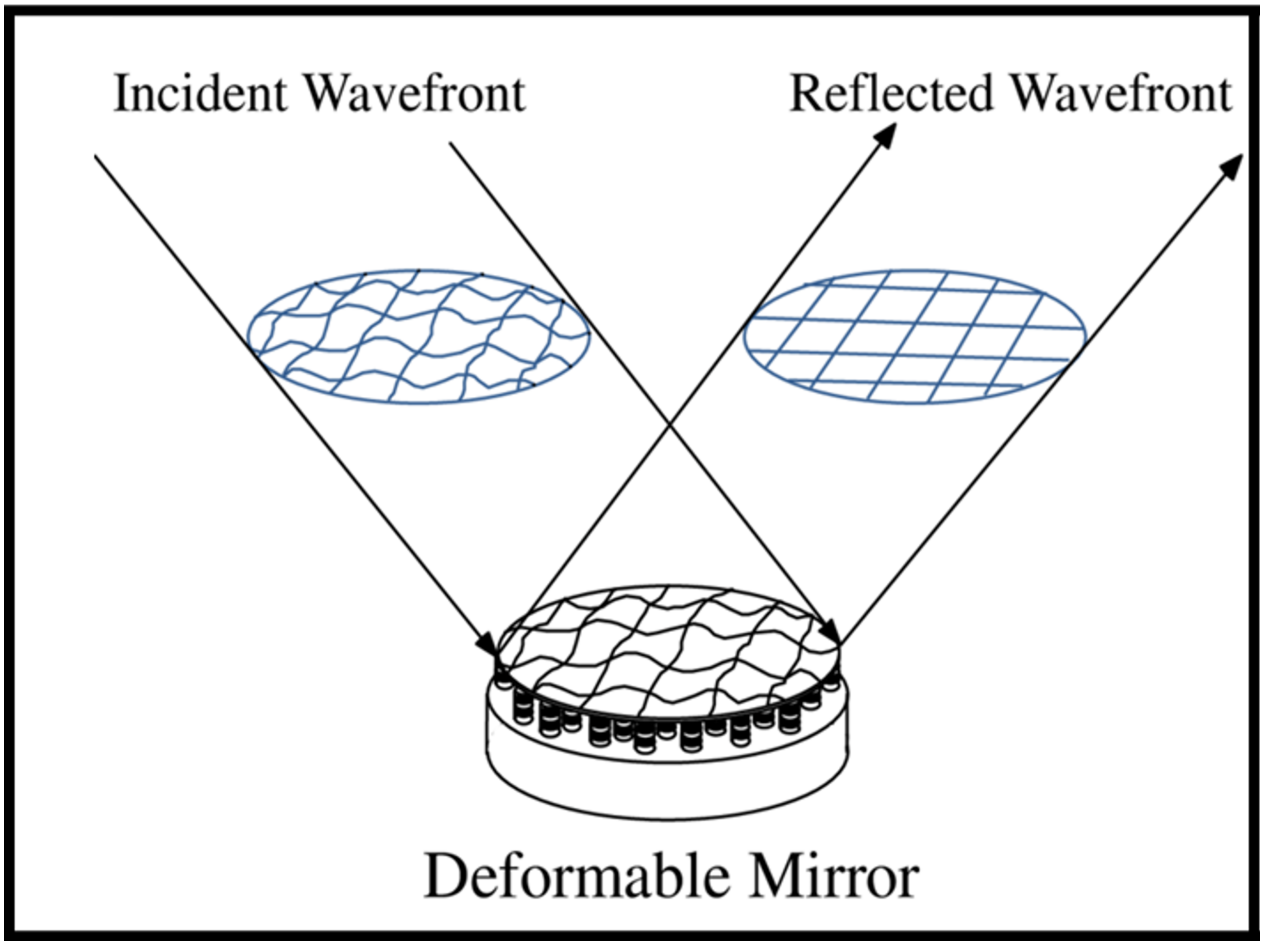

FIGURE 4.2 : Working principle of a DM.

4.2.2 Types of actuators used in deformable mirrors

Actuators are the essential components that enable shape modulation of the DM. Among the different types, electrostatic actuators (Zamkotsian et al., 2006) are widely used due to their high resolution, fast response time, and low power consumption. These actuators operate based on electrostatic attraction between a movable membrane and a fixed electrode, allowing for fine control over mirror deformation. However, their performance is constrained by a limited stroke range, which restricts the maximum displacement, and their susceptibility to voltage breakdown, which can affect long-term reliability. Despite these limitations, electrostatic actuators are particularly well-suited for applications requiring high-precision wavefront correction, such as adaptive optics in astronomy and compact optical systems. Their ability to provide rapid and highly accurate deformations makes them a key component in modern optical engineering.

Magnetic actuators (Cugat et al., 2001; Fernández et al., 2006) utilize the interaction between current-carrying conductors and a magnetic field to generate displacement, making them a valuable technology for DMs requiring large and precise shape adjustments. One of their key advantages is their large stroke range, which allows them to correct significant optical aberrations that

other actuator types may struggle to address. Additionally, they exhibit a highly linear response and low hysteresis, ensuring consistent and predictable performance. However, these benefits come at the cost of a relatively slower response time compared to electrostatic or piezoelectric actuators, as well as a bulkier design that demands more physical space within optical systems. Despite these limitations, magnetic actuators are particularly well-suited for applications such as correcting atmospheric distortions in ground-based telescopes, where maintaining high image quality over varying observing conditions is critical.

Piezoelectric actuators (Kanno et al., 2007; Toporovskiy et al., 2022) are a widely utilized technology in DMs due to their ability to generate substantial force and induce significant mirror deformations with high precision. These actuators operate based on the piezoelectric effect, where certain materials change shape in response to an applied electric field (or tension), exerting force on the mirror membrane. Their key advantages include high force output, fast response times in the sub-millisecond range, and excellent repeatability, making them suitable for dynamic wavefront control. However, their performance can be limited by hysteresis effects, which necessitate compensation strategies, and prolonged operation may induce material fatigue. Despite these constraints, piezoelectric actuators remain widely used in high-power laser systems due to their fast response and high precision. In this context, a piezoelectric bimorph DM equipped with 48 actuators, each operable within a voltage range of [–200 V, +200 V], was integrated into the *ALLS* 150 TW beamline to enable fine wavefront control and optimization.

4.2.3 Traditional bimorph deformable mirror

The design of a conventional bimorph DM has been presented in Figure 4.1. It consists of a passive substrate made of glass, silica, or copper and an active piezodisk firmly glued to the substrate. There is one “ground” electrode between the substrate and the piezodisk and a set of electrodes (actuators) placed on the outer surface of the piezodisk to reproduce various aberrations (Aleksandrov et al., 2004). When some voltage is applied to the control electrodes, the transverse piezoelectrical effect provides the mirror bend. If we consider, for example, a piezoceramic disk with surface S, then the change of surface ΔS for this shape is defined as :

Equation 4.1 : Surface deformation of a bimorph piezoelectric disk

$$\Delta S = \frac{2VSd_{13}}{t} \tag{4.1}$$

where d_{13} is the transverse piezoelectric coefficient quantifying how effectively the piezoceramic disk translates a vertical electric field into horizontal mechanical strain. Thus, it tells you how much radial strain (e.g., lateral deformation or displacement) results from an electric field applied across the thickness of the material. It usually comes from manufacturer datasheets obtained via defined

standards on piezoelectricity (Cain et al. (2014)). V is the applied voltage, and t is the thickness of the piezoceramic disk.

Bimorph DMs are particularly advantageous in laser-plasma acceleration experiments due to their high reliability, ease of integration, and ability to achieve large stroke deformations of the mirror surface. Bimorph DMs remain functional even in the event of individual electrode failure and deform in response to voltages applied to large, continuous electrodes that span broad regions of the mirror surface. Each electrode induces a global, smooth deformation, which resembles a specific optical aberration or mode (Roddier, 1988). This contrasts with "zonal" control (like in MEMS or stacked-actuator DMs), where each actuator affects a small, localized area independently. This makes bimorph DMs well-suited for compensating low-order wavefront aberrations, which is crucial for optimizing laser intensity distribution and focal spot quality in high-power laser systems (Lee et al., 2006; Samarkin et al., 2015). However, their primary limitation lies in their comparatively low spatial resolution, which restricts their ability to correct high-order aberrations in the wavefront. Addressing this issue requires increasing the density of control electrodes on the piezoceramic surface while minimizing insulation gaps, a trade-off that results in smaller electrode sizes and a corresponding decrease in the peak-to-valley deformation of each actuator. Maintaining an adequate local stroke necessitates reducing the mirror substrate thickness, which, in turn, increases the complexity of manufacturing, particularly in achieving high-quality mirror surface polishing. Despite these challenges, bimorph DMs remain a promising tool for enhancing laser-plasma interaction control, thereby contributing to more stable and efficient ion acceleration.

The curvature of a DM surface depends on different parameters and is given by the following equation (Ellis, 1999) :

Equation 4.2 : Radius of curvature of a bimorph deformable mirror

$$\rho_{\text{semi}} = \frac{1}{R} = \left(\frac{d_{13}E_{\max}}{t_1}\right)\frac{6kr(1+r)}{1+k^2r^4+2kr(2+3r+2r^2)} \tag{4.2}$$

where ρ_{semi} is the radius of curvature, $E_{\max}$ is the electric field, and $r = t_2/t_1$, t_1 and t_2 are the thicknesses of the piezoceramic plate and mirror substrate, respectively, and $k = E_2/E_1$, E_1 and E_2 are the Young's moduli of the piezoceramic material and substrate material, respectively. Here, the Poisson's coefficient (a material's property that describes the ratio of transverse strain to axial strain when a material is stretched or compressed) is assumed to be the same for both materials and therefore does not appear in Equation 4.2. The curvature of a bimorph DM arises from the strain mismatch between the bonded piezoceramic and substrate layers when an electric field is applied. The term $\frac{d_{13}E_{\max}}{t_1}$ quantifies the strain induced per unit thickness in the active layer, directly contributing to the overall bending. The fractional expression modulating this term captures the mechanical coupling governed by the relative thickness ($r = \frac{t_2}{t_1}$) and stiffness ($k = \frac{E_2}{E_1}$) of the two materials. Maximum curvature is obtained by minimizing the piezo layer thickness

($t_1 \downarrow$), maximizing the piezoelectric coefficient (d_{13}) and electric field ($E_{\max}$), and optimizing the mechanical ratios r and k.

The sensitivity S_b of the round bimorph DM could be expressed as the ratio between applied voltage and stroke of the mirror surface as :

Equation 4.3 : Sensitivity of a round bimorph deformable mirror

$$S_b = \frac{d^2}{8RV} = \frac{d^2}{4t_2^2} d_{13} \tag{4.3}$$

where d is the diameter of the piezoceramic disk. The first expression for sensitivity reflects the relationship between mirror stroke and curvature, with stroke proportional to $\frac{d^2}{8R}$, a standard result for spherical cap deformation. Substituting the radius of curvature R, derived from the piezoelectric bending model, leads to a simplified second form involving only the substrate thickness t_2 and the piezoelectric coefficient d_{13}, under the assumption of a strong mechanical asymmetry between layers. A larger mirror diameter or a thinner substrate results in greater sensitivity. Sensitivity improves with a high d_{13} and a small t_2, aligning with more compliant mirror designs. Thin piezoelectric layers combined with stiff substrates enable larger curvatures and enhance the possibilities for wavefront shaping. The sensitivity S_b is essential for selecting the actuator voltage range needed to achieve a desired stroke, making this framework critical in the design and calibration of DMs for adaptive optics applications.

In summary, real-time wavefront correction depends on the tight integration of precise sensing and reliable modulation. At our facility, this modulation is provided by a traditional bimorph DM, which adjusts the optical wavefront via voltage-controlled surface deformations. Its performance, however, hinges on real-time feedback from a high-resolution WFS (a *SID4* WFS from *Phasics*), completing the formation of a closed-loop system enabling wavefront shaping.

4.2.4 *SID4* wavefront sensor from *Phasics* : Real-time feedback in adaptive optics

The *SID4* WFS, developed by *Phasics*, is a high-resolution Quadriwave Lateral Shearing Interferometry (QWLSI)-based WFS designed to measure and analyze optical wavefronts in real time (Chanteloup, 2005). This method whose principle is shown in Figure 4.3, operates as follows :

1. **Diffractive grating encoding :** The incoming beam passes through a specialized diffraction grating (MHM), creating multiple replicas of the wavefront.
2. **Interference Pattern Formation :** These replicas interfere at a fixed lateral shear on a detection plane, generating an interferogram recorded by a high-resolution CCD or CMOS sensor.
3. **Phase gradient measurement :** By analyzing the spatial phase gradients from the interferogram, the wavefront is reconstructed.

4. **Wavefront reconstruction :** Using proprietary reconstruction algorithms, the absolute wavefront phase map is obtained, allowing detailed characterization of optical aberrations in both Zernike and Legendre formalisms.

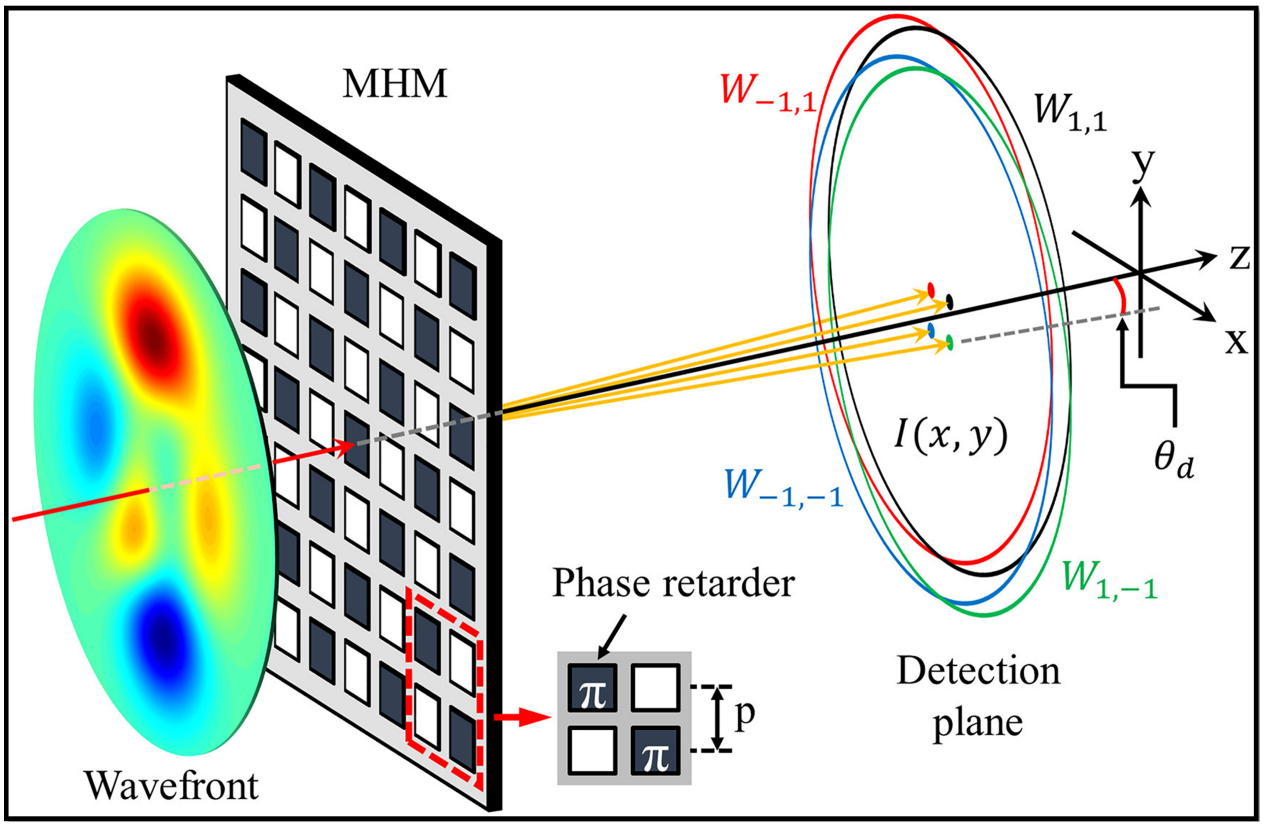

FIGURE 4.3 : Quadriwave Lateral Shearing Interferometry.

This figure has been adapted from Song et al. (2024), with the publisher's permission.

The wavefront phase $W(x, y)$ is obtained by integrating the measured phase gradients :

Equation 4.4 : Wavefront phase reconstruction

$$W(x, y) = \int \nabla W \, \mathrm{d}x \, \mathrm{d}y \tag{4.4}$$

where $W(x, y)$ represents the wavefront phase, and ∇W is derived from the measured intensity gradients.

At the *ALLS* 150TW facility, the *SID4* wavefront sensor (WFS) based on quadriwave lateral shearing interferometry (QWLSI) is coupled with a bimorph deformable mirror (DM) in a closed-loop configuration to correct optical aberrations. The real-time feedback loop operates as follows :

1. **Initial wavefront measurement :** The WFS captures the uncorrected wavefront profile $\phi_0(x, y)$ of the laser beam. The RMS wavefront phase error is computed as

$$\mathrm{RMS} = \left(\frac{1}{|\Omega|} \int_{\Omega} [\phi_0(x, y) - \phi_{\text{ref}}(x, y)]^2 \, \mathrm{d}A \right)^{1/2},$$

   where ϕ_{ref} represents the desired (typically planar) wavefront and Ω the sensor aperture.

2. **DM calibration :** The calibration process establishes the influence matrix $\mathbf{M}$ that links actuator voltages to measured phase changes. A known voltage is sequentially applied to each actuator, and the resulting phase deformation is recorded by the WFS. Assuming a linear mirror response, $\mathbf{M}$ can later be inverted to compute the voltages required to produce a specific corrective phase map.

3. **Actuator adjustment and response measurement :** The measured wavefront is decomposed into a modal basis, typically a set of Zernike polynomials $Z_j(x, y)$:

$$\phi_k(x, y) \approx \sum_{j=1}^{J} a_j^{(k)} Z_j(x, y),$$

   where $a_j^{(k)}$ are the modal coefficients at iteration k. Low-order aberrations (tilt, defocus, astigmatism, coma, spherical) are compensated first. The DM voltages are updated to minimize the RMS error according to a regularized least-squares criterion :

$$\Delta \mathbf{V} = (\mathbf{M}^\top \mathbf{M} + \lambda^2 \mathbf{I})^{-1} \mathbf{M}^\top (\boldsymbol{\phi}_{\text{ref}} - \boldsymbol{\phi}_k),$$

   followed by a conservative update $\mathbf{V}_{k+1} = \mathbf{V}_k + g \, \Delta \mathbf{V}$ with a gain factor $0 < g \leq 1$.

4. **Instantaneous correction validation :** After each voltage update, the new wavefront ϕ_{k+1} is measured, and the residual RMS error is evaluated. The loop continues until the RMS convergence criterion $\Delta \mathrm{RMS} < \epsilon$ is reached or the target RMS value is achieved.

5. **Optimized configuration determination :** The final actuator voltage vector corresponds to the configuration that minimizes aberrations according to the RMS wavefront phase error.

In practice, convergence is typically achieved within a few iterations, yielding a near-diffraction-limited beam quality.

The quantitative impact assessment of different DM actuator configurations relies on wavefront quality metrics, such as :

— **Peak-to-Valley (PtV) wavefront error :** The maximum optical path difference between the highest and lowest points of the wavefront. It is sensitive to localized aberrations but does not provide a statistical measure of overall wavefront quality.

— **RMS wavefront error :** The square root of the spatial average of the squared deviations from a reference wavefront, often a plane. It offers a robust statistical indicator of the overall wavefront distortion and is less affected by isolated defects than PtV.

— **Strehl ratio :** The ratio of the peak intensity of the actual point spread function (PSF) to that of an ideal, aberration-free system. It provides a normalized measure (ranging from 0 to 1) of optical quality, with values above 0.8 typically considered diffraction-limited.

The Strehl ratio can be approximated for values above 0.3–0.4, as expressed by Mahajan (1983) :

Equation 4.5 : Strehl ratio in the Fraunhofer approximation

$$\text{Strehl Ratio} = \exp\Big[-(2\pi\sigma_\phi)^2\Big] \tag{4.5}$$

where σ_ϕ is the RMS wavefront phase error.

At the *ALLS* 150 TW facility, the *SID4* wavefront sensor plays a pivotal role in the real-time characterization and optimization of the laser beam wavefront before target interaction. Its high-resolution phase retrieval enables the accurate detection and correction of aberrations induced by optical components or thermal effects, and supports precise wavefront tailoring through DM actuation. In contrast to traditional Shack–Hartmann sensors or Fizeau interferometers (see Figure 4.4), the *SID4* achieves a nanometric phase resolution < 2 nm RMS and an absolute accuracy of $\sim$20 nm RMS over a $\sim$10 mm pupil, as specified by the manufacturer. The dynamic measurement range reaches up to ~ 500 μm PtV, which enables simultaneous detection of small residual aberrations and large-scale distortions within a single acquisition. By comparison, high-quality Shack–Hartmann sensors typically offer phase resolutions of 10–20 nm RMS with dynamic ranges below 50–100μm PtV, limited by microlens pitch and spot centroiding accuracy, while Fizeau interferometers, although capable of $\lambda/200$ ($\simeq$4 nm at 800 nm) accuracy, require vibration isolation and are less suited for single-shot or high-power environments. The *SID4* thus combines the nanometric precision of interferometry with the robustness and temporal resolution of imaging sensors, making it ideally suited for closed-loop feedback control in adaptive wavefront shaping experiments.

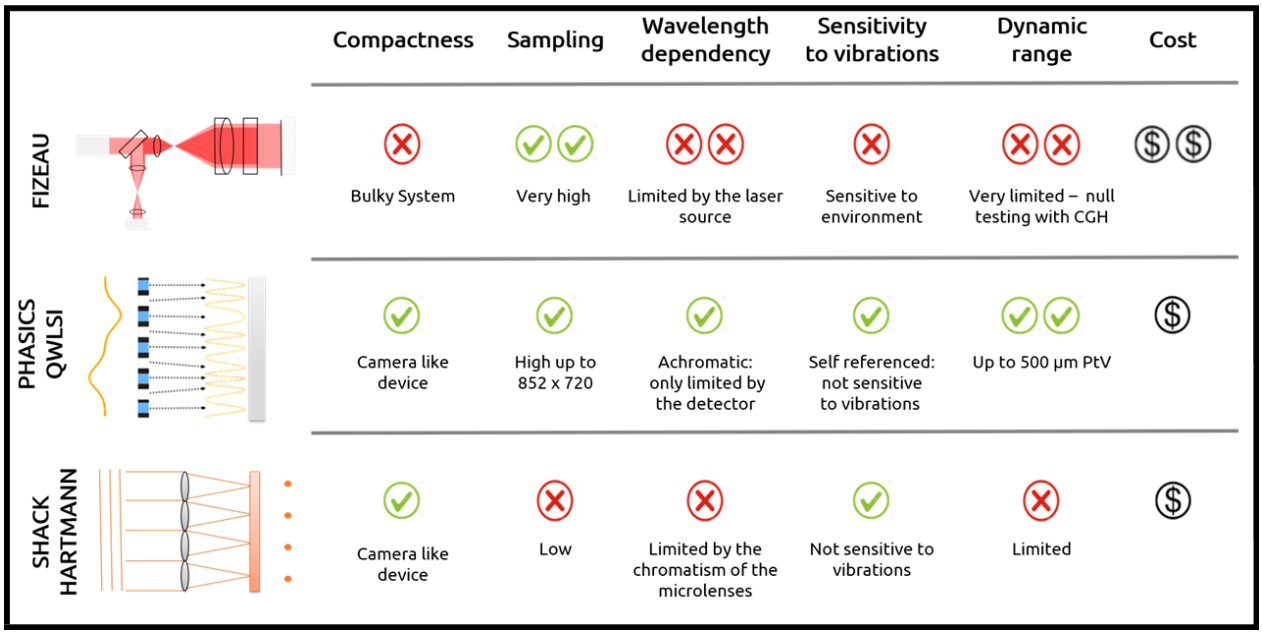

FIGURE 4.4 : Comparison of wavefront sensing techniques.

This figure has been adapted from Phasics (2025), with the publisher's permission.

4.3 Pulse characterization and compression

Precise characterization of ultrashort laser pulses is critical for understanding and controlling light-matter interactions in high-intensity regimes. Conventional techniques such as FROG (Trebino et al., 1997) and SPIDER (Iaconis et al., 1998) suffer from limitations in temporal resolution and require complex phase retrieval algorithms. On the *ALLS* 150 TW beamline, optimal pulse characterization is achieved by using a *Wizzler* (Amplitude Laser, 2025b) device developed by *Fastlite* to measure the spectral phase of the laser pulse and dynamically adjust it via feedback to a *Dazzler* AOPDF developed by *Fastlite*. The *Dazzler* is inserted between the two CPA stages as a spectral-phase and dispersion management tool. It's not for compression itself, but for programmable pre-compensation and shaping before the final amplification and compression. Thus, the *Dazzler* allows dynamic adjustment of dispersion (β_2, β_3, ...) and spectral phase shaping. This is critical for :

— Optimizing pulse duration after final compression.

— Pre-compensating for known material dispersion in downstream optics (mirrors, compressor gratings, target chamber windows).

— Tailoring the temporal profile of the pulse.

The *Wizzler* is an advanced pulse measurement system specifically designed for precise spectral phase characterization of amplified ultra-short, near-Fourier-transform-limited laser pulses. It utilizes Self-Referenced Spectral Interferometry (SRSI), a patented technique developed by *FAST-*

LITE. It is single-beam, single-shot, and delivers both spectral phase and amplitude measurements, i.e, the complete temporal characterization of an ultrashort pulse.

An ultrashort laser pulse in the spectral domain is given by :

Equation 4.6 : Spectral representation of an ultrashort laser pulse

$$E(\omega) = |E(\omega)|\, e^{i\phi(\omega)} \tag{4.6}$$

where $|E(\omega)|$ is the spectral amplitude and $\phi(\omega)$ is the spectral phase.

The spectral phase can be expanded as :

Equation 4.7 : Spectral phase expansion around the central frequency

$$\phi(\omega) = \phi_0 + \tau_g(\omega - \omega_0) + \frac{\beta_2}{2}(\omega - \omega_0)^2 + \frac{\beta_3}{6}(\omega - \omega_0)^3 + \dots \tag{4.7}$$

where :

— τ_g is the group delay describing how different frequency components travel through the optical system. If all frequency components are delayed by the same τ_g, the pulse maintains its shape. Only its arrival time shifts.

— $\beta_2 = \frac{\mathrm{d}\tau_g}{\mathrm{d}\omega}$ is the GDD and describes how fast the group delay changes with frequency, introducing pulse broadening (linear chirp),

— $\beta_3 = \frac{\mathrm{d}^2\tau_g}{\mathrm{d}\omega^2}$ is the TOD which causes asymmetric pulse distortions.

SRSI enables direct retrieval of the spectral phase of an ultrashort pulse by utilizing Cross-Polarized Wave (XPW) generation as a self-referenced phase measurement. The resulting interferogram encodes phase-dependent modulations :

Equation 4.8 : Interferometric spectral intensity retrieved by SRSI

$$I(\omega) = |E(\omega)|^2 + |E_{\mathrm{XPW}}(\omega)|^2 + 2\,|E(\omega)|\,|E_{\mathrm{XPW}}(\omega)|\cos(\Delta\phi(\omega)), \tag{4.8}$$

where $E_{XPW}(\omega)$ is the XPW-generated reference pulse and $\Delta\phi(\omega) = \phi(\omega) - \phi_{XPW}(\omega)$ is the relative phase.

Applying a Fourier-domain filtering technique allows direct retrieval of the spectral phase $\phi(\omega)$, eliminating the need for iterative phase retrieval algorithms such as those used in SPIDER or FROG.

4.4 Working mechanism of the *Wizzler*

4.4.1 Temporal pulse reconstruction

To measure the duration of ultrashort laser pulses with a *Wizzler* device, a fraction of the input beam generates a reference pulse through XPW generation in a nonlinear crystal. The XPW pulse has a broader spectrum and a flatter spectral phase than the original pulse. In parallel, a delayed replica of the input pulse is created using a birefringent plate. The XPW pulse and the replica interfere at the spectrometer, producing a spectral interferogram. By applying Fourier-Transform Spectral Interferometry (FTSI) (Lepetit et al., 1995; Dorrer et al., 2001) analysis, the *Wizzler* retrieves both the spectral phase and spectral amplitude of the input pulse, and reconstructs its temporal profile. From the extracted spectral phase, dispersion coefficients such as β_2, β_3, and higher orders can be precisely determined.

4.4.2 Experimental implementation

The *Wizzler* is typically integrated into the laser chain as follows :

1. A small fraction of the laser beam is picked off and directed to the *Wizzler* diagnostic line.
2. The device retrieves the spectral phase and reconstructs the corresponding temporal profile, including the pulse duration.
3. These measurements are then used to refine and apply the desired spectral phase corrections through the *Dazzler*.

4.4.3 Spectral phase compensation using the *Dazzler*

The *Dazzler* is used to shape the spectral phase of the *ALLS* laser. The core mechanism relies on acousto-optic interaction in a birefringent crystal. A radio-frequency (RF) signal, modulated in time, generates a traveling acoustic wave along the crystal's z-axis. This acoustic wave forms a dynamic volume grating whose local spatial frequency varies along the z-axis and encodes the RF signal shape. When a chirped optical pulse (initially in the ordinary polarization) enters the crystal, each optical frequency ω encounters a specific grating region $z(\omega)$ where phase-matching is satisfied. At that location, part of the light is diffracted into the extraordinary polarization. This frequency-selective diffraction causes each spectral component to be delayed differently because the group velocities of the ordinary and extraordinary polarizations are not the same. The result is a programmable spectral phase and amplitude modulation of the optical pulse. By tailoring the RF waveform, the *Dazzler* allows for precise control over GDD, TOD, and even higher-order spec-

tral phase terms. Figure 4.5 illustrates the interaction between a chirped acoustic wave and an ultrashort optical pulse, highlighting the frequency-dependent diffraction and induced group delay.

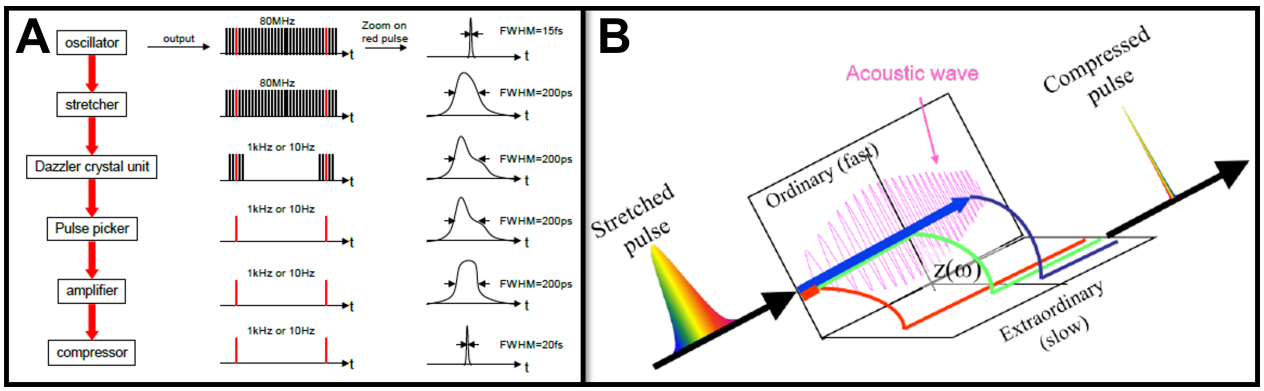

FIGURE 4.5 : *Dazzler* principle in CPA laser.

Setup of an ultrafast CPA laser with a *Dazzler*. (B) Explanation of the *Dazzler* principle in terms of group delay control. These figures have been adapted from the *Dazzler* user manual : Amplitude Laser (2019), with the publisher's permission.

The group delay τ_g applied to the diffracted pulse can be expressed by Tournois (1997) :

Equation 4.9 : Group delay induced by the *Dazzler* crystal

$$\tau_g(\omega) = \frac{n_{g_o}(\omega)}{c}\, z(\omega) + \frac{n_{g_e}(\omega)}{c}\, (L - z(\omega)), \tag{4.9}$$

where n_{g_o} and n_{g_e} are respectively the ordinary and extraordinary group indexes along the propagation direction, c the speed of light and L the crystal length.

Controlling for each optical frequency ω, the position $z(\omega)$ where ω is diffracted enables control of the pulse group delay. The amplitude of the output pulse is controlled by the acoustic power at position $z(\omega)$.

4.5 Experimental setup to perform Target Normal Sheath Acceleration

Experiments were conducted on the laser-driven ion acceleration beamline at the *Institut national de la recherche scientifique* (*INRS*) in Varennes, Canada, using the 150 TW Ti : Sapphire laser system at the *ALLS* facility (Vallières et al., 2020). The experimental setup is depicted in Figure 4.6. This dual-stage CPA system delivers 3.2 J on target with a minimum pulse duration of 22 fs at 800nm. Temporal shaping was achieved with a *Dazzler* AOPDF, enabling independent tuning of GDD and TOD, and a closed-loop *Wizzler–Dazzler* system was first used for optimization through dispersion control of the pulse duration from 32 fs to 22 fs. In our setup, the *Wizzler* is positioned on a diagnostic arm that samples a small fraction of the main beam after the compres-

sor, and therefore after both the DM and the *Dazzler*. This configuration ensures that the spectral phase characterization includes all dispersive and corrective elements acting on the beam before focusing and target interaction, providing an accurate measurement of the pulse as delivered to the experimental area. Downstream, spatial wavefront shaping is performed using a piezoelectric bimorph DM (Figure 4.7) from *AKAOptics*, equipped with 48 actuators operating over a [-200 V, +200 V] range. The laser beam's spatial profile is optimized through a closed-loop configuration that couples the DM with a *SID4* WFS from *Phasics*, iteratively minimizing the RMS wavefront error to achieve a near-flat-wavefront spot. Combined with the shortest pulse duration (22 fs), it defined the reference setup from which independent variations in spatial and temporal shaping of TNSA were explored. Finally, the beam is directed toward the target and tightly focused by an $f/3$ off-axis parabola to a focal spot size of approximately 5 μm, reaching a peak intensity of $I_0 \sim 1.3 \times 10^{20}$ W.cm^{-2}. To ensure good temporal contrast, the laser pulse undergoes XPW generation before final compression, resulting in an amplified spontaneous emission contrast better than 10^{-10} at -100 ps before the main pulse, and a steep power rise with contrast below 10^{-6} at -3 ps.

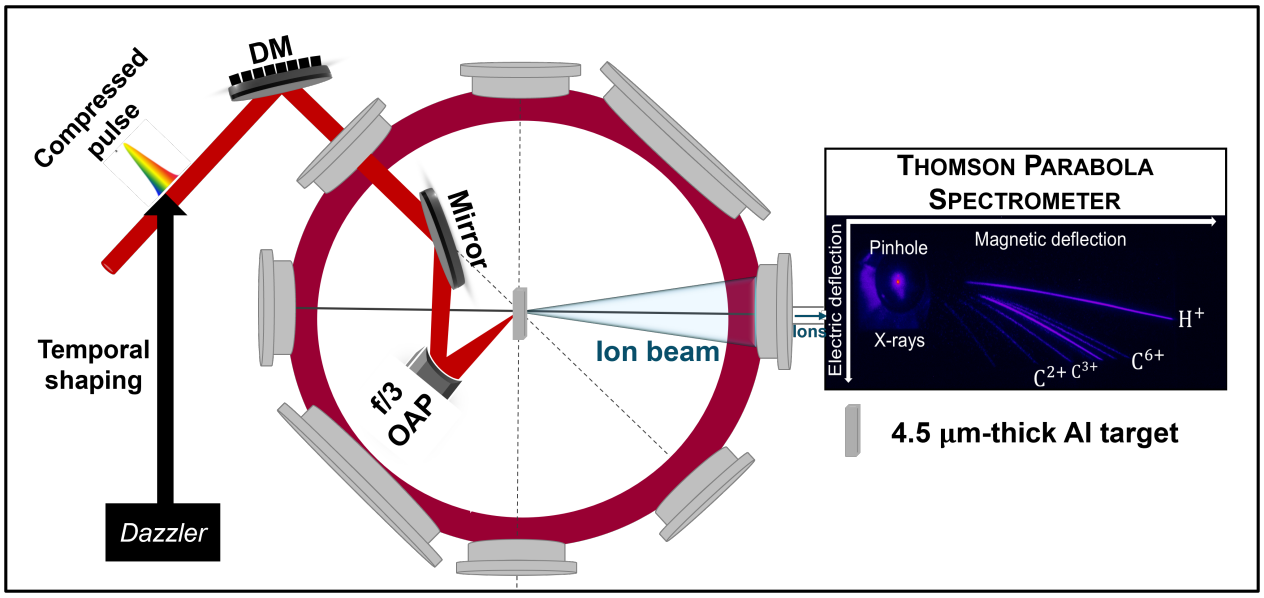

FIGURE 4.6 : Experimental setup layout.

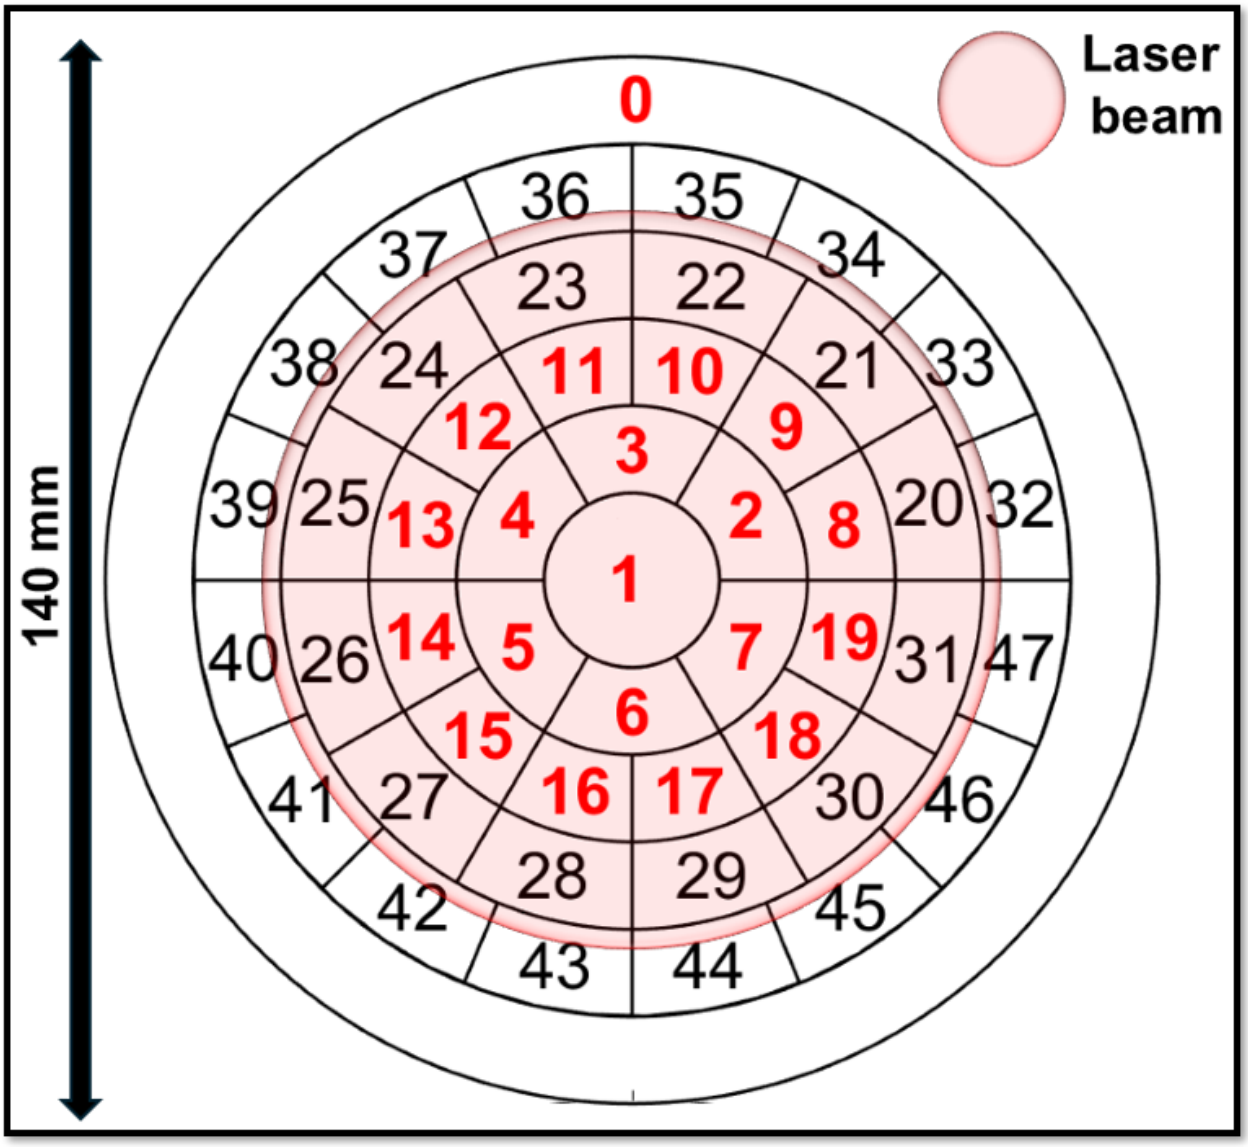

FIGURE 4.7 : Layout of the 48 actuators on the DM.

The 48 actuators on the deformable mirror are used for laser wavefront control. Actuators are numbered from 0 to 47, with those utilized for the analysis highlighted in red (actuators 0–19). Actuator 0 corresponds to the global defocus mode, which translates the entire mirror surface forward or backward.

The experiment employed p-polarized laser pulses incident at an angle of 20° relative to the target normal to irradiate 4.5-μm-thick aluminum foils, thereby generating ion beams via the TNSA mechanism. Targets were mounted on a multi-target holder capable of accommodating up to 400 targets, allowing for high-throughput operation at a repetition rate of 0.625 Hz (Catrix et al., 2023). The target was precisely positioned at the laser focus to ensure consistent interaction conditions. Ion spectra were measured shot-by-shot using a Thomson parabola spectrometer (TPS) coupled with a microchannel plate (MCP) detector. Additional details on the diagnostic configuration and experimental setup can be found in Catrix et al. (2023) and Vallières et al. (2020).

4.6 Methodology : Proton beam characterization and adaptive laser parameter tuning

The properties of laser-driven proton beams, such as energy distribution, yield, and spectral shape, are sensitive to the laser spatio-temporal aspect at the time of interaction, making precise parameter optimization essential for improving proton beam quality and stability. For each laser shot, key performance metrics are extracted from the ion spectra recorded by the TPS using a custom *Matlab* code (see Chapter 3).

The total proton yield per steradian (N_{tot} [protons.sr^{-1}]) is obtained by integrating the proton yield $\frac{d^2N}{d\mathcal{E}d\Omega}$ [protons.MeV^{-1}.sr^{-1}] over energy :

Equation 4.10 : Energy-integrated proton yield

$$N_{\text{tot}} = \int \frac{d^2N}{d\mathcal{E}\, d\Omega}\, d\mathcal{E}, \tag{4.10}$$

with $\mathcal{E}$ [MeV] the proton energy, N [protons] the total number of protons and Ω the solid angle occupied by the laser-driven proton beam.

The total proton energy, $\mathcal{E}_{tot}$ [MeV], is computed by weighting the proton distribution by energy, providing an estimate of the overall energy carried by the beam (Snavely et al., 2000; Wilks et al., 2001) :

Equation 4.11 : Total proton beam energy

$$\mathcal{E}_{\text{tot}} = \iint \frac{d^2N}{d\mathcal{E}\, d\Omega}\, \mathcal{E}\, d\mathcal{E}\, d\Omega \tag{4.11}$$

The maximum proton energy (cut-off energy), $\mathcal{E}_{p,max}$ [MeV], typically defined experimentally as the proton energy located at the midpoint of the final slope in the spectrum, serves as a crucial parameter in assessing the efficiency of the acceleration mechanism (Fuchs et al., 2006; Robson et al., 2007; Flippo et al., 2008; Zeil et al., 2010).

The mean proton energy, $\bar{\mathcal{E}}$ [MeV], characterizes the average energy of the proton beam and is given by :

Equation 4.12 : Mean proton energy

$$\bar{\mathcal{E}} = \frac{\iint \frac{d^2N}{d\mathcal{E}\, d\Omega}\, \mathcal{E}\, d\mathcal{E}\, d\Omega}{N_{\mathrm{tot}}} \tag{4.12}$$

Similarly, the second-order moment in energy ($\bar{\mathcal{E}^2}$ [MeV^2]) is defined as :

Equation 4.13 : Second-order moment in energy

$$\bar{\mathcal{E}^2} = \frac{\iint \frac{d^2N}{d\mathcal{E}\, d\Omega}\, \mathcal{E}^2\, d\mathcal{E}\, d\Omega}{N_{\mathrm{tot}}} \tag{4.13}$$

Both $\bar{\mathcal{E}}$ and $\bar{\mathcal{E}^2}$ play a significant role in defining the energy distribution width and spectral shape.

Additionally, the hot electron temperature T_{hot}, inferred from the slope of the proton spectrum on a logarithmic scale, provides a crucial link between laser-plasma interaction physics and the acceleration process (Mora, 2003; Fuchs et al., 2006; Passoni et al., 2010).

These metrics collectively enable a comprehensive analysis of proton beam properties, guiding optimization efforts in TNSA experiments.

To enable high-repetition-rate laser-driven ion acceleration at 0.625 Hz while actively tuning laser parameters, an adaptive feedback system has been implemented to adjust the voltages of the DM dynamically. The DM control system was developed in *MATLAB* using the *SID4 OASys* software development toolkit from *Phasics*. Before each laser shot, a specific DM voltage configuration is applied, shaping the wavefront before the laser pulse interacts with the target. To facilitate the configuration and deployment of DM settings, a custom *MATLAB* GUI was developed. This GUI enables users to modify actuator voltages, load mirror drivers, monitor high-voltage status, and generate voltage maps. It also provides options to define actuator ranges, voltage limits, and step sizes, offering experimental flexibility. Simultaneously, the *Dazzler* is remotely controlled via a servor to perform shot-to-shot spectral phase shaping at the same 0.625 Hz repetition rate, allowing independent tuning of GDD and TOD.

4.6.1 Dispersion-induced variations in proton beam metrics

In the context of ultrafast optics, the electric field of an initially Fourier-limited Gaussian pulse in the frequency domain is given by :

Equation 4.14 : Electric field of a Fourier-limited Gaussian pulse in the spectral domain

$$\tilde{E}(\omega) = E_0 \, \exp\left[-\frac{(\omega - \omega_0)^2}{2\Delta\omega_0^2}\right] \tag{4.14}$$

where ω_0 is the central angular frequency and $\Delta\omega_0$ is the spectral width (related to the FWHM pulse duration via $\tau_0 = \frac{2\sqrt{\ln 2}}{\Delta\omega_0}$). When a spectral phase $\phi(\omega)$ is applied, such as through a dispersive medium or an AOPDF, the modified spectrum becomes :

Equation 4.15 : Modified electric field spectrum through a dispersive medium

$$\tilde{E}_{\text{out}}(\omega) = \tilde{E}(\omega) \, \exp[i\phi(\omega)] \tag{4.15}$$

where the spectral phase is typically expanded in a Taylor series around ω_0 and dominated by its second- and third-order terms. In the weak-dispersion regime, the output pulse duration τ_{out} can be estimated using the second-moment method :

Equation 4.16 : Output pulse duration after dispersion through a *Dazzler*

$$\tau_{\text{out}} = \tau_0 \sqrt{1 + \left(\frac{4\ln(2)\,\beta_2}{\tau_0^2}\right)^2 + \left(\frac{4(\ln(2))^{3/2}\,\beta_3}{\tau_0^3}\right)^2} \tag{4.16}$$

where :

— τ_0= 22 fs : Initial transform-limited pulse duration (FWHM), i.e., the shortest possible duration with minimal spectral phase distortion.

— β_2 : GDD applied by the *Dazzler*, in fs^2.

— β_3 : TOD applied by the *Dazzler*, in fs^3.

— $4\ln(2)$, $(\ln(2))^{3/2}$: Constants arising from the relation between FWHM and RMS duration for a Gaussian temporal profile.

This formula is derived under the assumption of small spectral phase distortions and remains valid in the weak dispersion regime, typically for $|\beta_2| \ll \tau_0^2$ and $|\beta_3| \ll \tau_0^3$. For each set of dispersion values β_2 and β_3 applied by the *Dazzler*, the resulting pulse duration τ_{out} of the modified pulse is determined by fitting a Gaussian function to its temporal intensity profile even if TOD introduces asymmetric distortions and create sub-structures or satellite pulses due to phase oscillations in the spectral phase. Figure 4.8 illustrates the temporal broadening of the *ALLS 150 TW* laser pulse across a range of β_2 and β_3 values due to the spectral phase imparted by the *Dazzler*. These values were used to guide the laser pulse shaping conditions applied in the experiment.

Although a Gaussian pulse shape is exact only for transform-limited pulses with symmetric temporal profiles, it remains a convenient and consistent metric to extract an effective duration τ_{out} from

experimental or simulated traces. In our analysis, the temporal intensity $I(t)$ is systematically fitted using a skewed-Gaussian model :

Equation 4.17 : Skewed-Gaussian temporal intensity profile

$$I(t) = I_0 \exp\left[-4\ln(2)\left(\frac{t-t_0}{\tau_{\text{out}}}\right)^2\right]\left[1+\operatorname{erf}\left(s\,\frac{t-t_0}{\tau_{\text{out}}}\right)\right] \tag{4.17}$$

where I_0 is the peak intensity, t_0 the temporal centroid, τ_{out} the full width at half maximum (FWHM) of the Gaussian envelope, and s a skew parameter. The fitted width τ_{out} is reported as the "pulse duration" for comparison between dispersion settings, while the skew parameter s captures temporal asymmetry. This two-parameter fit provides a robust description of the main pulse even in the presence of higher-order spectral phase, such as TOD, which can generate pre-pulses, post-pulses, or asymmetric wings.

Figure 4.9 illustrates representative temporal intensity profiles obtained for different β_2 and β_3 values applied by the *Dazzler*. The extracted parameters τ_{out} and s allow direct visualization of both the broadening and asymmetry trends across the dispersion scan.

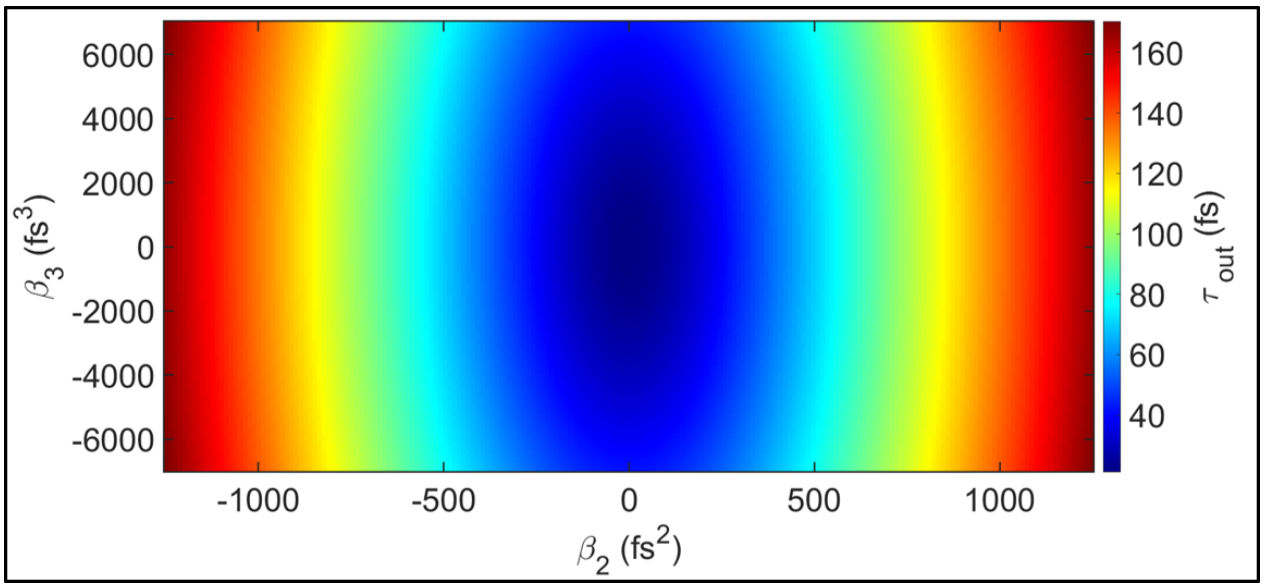

FIGURE 4.8 : Pulse duration vs β_2 and β_3.

Temporal broadening, or increasing pulse duration, of *ALLS* 150 TW laser pulse after travelling through the *Dazzler* crystal, plotted as a function of β_2 and β_3.

In this experiment, we performed a sensitivity analysis of proton beam metrics as a function of the laser pulse's spectral phase, specifically the β_2 and β_3. By systematically varying β_2 within the range [-1250,1250] fs^2 and β_3 from -7000 to 7000 fs^3 using the *Dazzler*, we assess the impact of temporal pulse shaping on proton beam generation. Prior to the experimental campaign, the

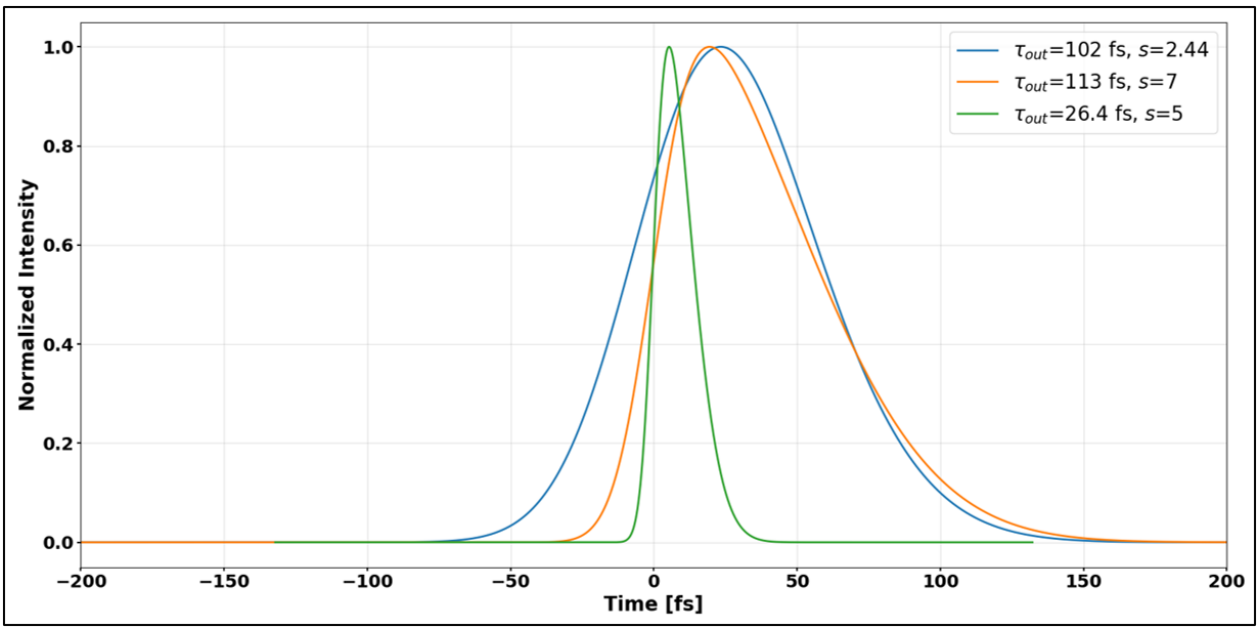

FIGURE 4.9 : Asymmetric Gaussian fits of pulse intensity profiles.

Intensity profiles fitted with an asymmetric Gaussian model, defined by an effective pulse duration τ_{out} and a skewness parameter s. The comparison illustrates how increasing s leads to stronger temporal asymmetry and a sharper rising edge, while larger τ_{out} values broaden the overall envelope.

range of *Dazzler* parameters (e.g., GDD and TOD extrema) to be explored was carefully identified to ensure machine safety and spectral integrity. Each configuration was pre-calculated to confirm that the applied phase would not introduce discontinuities or spectral holes in the laser pulse spectrum. This preventive step ensured that only physically consistent and optically safe settings were uploaded to the *Dazzler* before firing high-power shots.

Moreover, to evaluate the influence of temporal pulse shaping on proton beam metrics in our experimental setup, two key statistical metrics are used to assess whether temporal pulse shaping via GDD and TOD is relevant for further TNSA optimization efforts. The mean shot-to-shot standard deviation, denoted as σ_{sts}, characterizes the average fluctuation in proton beam metrics across repeated laser shots, providing a measure of the system's intrinsic temporal stability. For each combination of β_2 and β_3, 3 consecutive measurements were acquired. The standard deviation was first calculated across these 3 shots for each dispersion setting referred to as σ_i. The mean variability over all configurations was then computed as :

Equation 4.18 : Mean shot-to-shot standard deviation of proton beam metrics

$$\sigma_{\mathrm{sts}} = \frac{1}{\sqrt{N}} \sqrt{\sum_{i=1}^{N} \sigma_i^2} \tag{4.18}$$

where :

— σ_i denotes the standard deviation calculated over 3 laser shots for the i-th β_2-β_3 combination,

— N=114 is the total number of tested β_2-β_3 combinations.

The second statistical metric quantifies the variability of mean proton beam metrics across all tested β_2–β_3 configurations. Denoted as σ_{config}, this metric reflects how much the average beam performance varies from one dispersion setting to another. It is defined as :

Equation 4.19 : Configuration-to-configuration variability of proton beam metrics

$$\sigma_{\mathrm{config}} = \frac{1}{\sqrt{N}} \sqrt{\sum_{i=1}^{N} (\bar{x}_{\mathrm{mean}} - \bar{x}_i)^2} \tag{4.19}$$

where :

— $\bar{x}_i$ is the mean value of the proton beam metric for the i-th β_2–β_3 combination (averaged over 3 shots),

— $\bar{x}_{\mathrm{mean}}$ is the overall average across all configurations,

— N=114 is the total number of tested β_2-β_3 combinations.

This metric provides a measure of the spread in average beam performance across configurations, offering insight into the sensitivity of the acceleration process to spectral phase shaping.

Table 4.1 shows that σ_{sts} is nearly identical or higher than σ_{config} regardless of the proton beam metric, it suggests that the variability within individual β_2-β_3 combinations is similar or higher to the variability across different β_2-β_3 combinations. This indicates that changing β_2 and β_3 does not introduce much additional variability beyond the inherent fluctuations in the system and that the spectral phase changes do not have a strong systematic effect on the TNSA efficiency across the explored β_2-β_3 parameter space.

TABLE 4.1 : Sensitivity of proton metrics to temporal dispersion.

Proton beam metrics	$\sigma_{\mathbf{sts}}$	$\sigma_{\mathbf{config}}$	$\sigma_{\mathbf{config}}/\sigma_{\mathbf{sts}}$
N_{tot} [10^{10} protons·sr^{-1}]	0.77	0.61	0.79
$\mathcal{E}_{tot}$ [J]	2.62×10^{-3}	2.31×10^{-3}	0.88
$\mathcal{E}_{p,max}$ [MeV]	0.32	0.35	1.09
$\bar{\mathcal{E}}$ [MeV]	8.06×10^{-2}	6.56×10^{-2}	0.81
$\bar{\mathcal{E}^2}$ [MeV^2]	0.64	0.65	1.02

Statistics on the sensitivity of proton beam metrics to temporal dispersion. Comparison of shot-to-shot standard deviation ($\sigma_{\mathbf{sts}}$), across-configuration variability ($\sigma_{\mathbf{config}}$), and their ratio for each metric as a function of GDD and TOD.

4.6.2 Sensitivity analysis of proton beam metrics concerning deformable mirror actuator voltages

First, a matrix of 360 DM configurations was generated by varying the voltages of actuators 0–19 within the range [-80 V, 80 V], using Latin Hypercube Sampling (LHS) (Shields et al., 2016; Helton et al., 2003) to ensure diverse and well-distributed coverage of the 20-dimensional parameter space. These variations were applied on top of the initial configuration provided by the *SID4* wavefront sensing software, serving as the baseline for further adjustments. Subsequently, systematic variations were introduced to the first 20 actuators, and each resulting configuration was tested experimentally to evaluate the correlation between DM voltages and key proton beam metrics.

The correlation matrix in Figure 4.10.A highlights that only a small subset of the 20 DM actuators significantly influence the main proton beam metrics. Among them, Actuator0 stands out, showing consistently strong positive correlations with values exceeding 0.8 for $\mathcal{E}_{p,max}$, $\bar{\mathcal{E}}$, $\bar{\mathcal{E}^2}$, $\mathcal{E}_{tot}$ and T_{hot}. Its absolute average correlation across all metrics (Figure 4.10.B) is substantially higher than that of the other actuators, indicating that it has a dominant effect on the effective intensity delivered on target. Actuator1, and to a lesser extent Actuator16 and Actuator3, also show moderate influence, while the remaining actuators have minimal impact. This suggests that the system is susceptible to a reduced set of actuators, enabling efficient optimization even within a lower-dimensional parameter space. Moreover, Figure 4.11 shows weak inter-actuator correlations (<|0.17|), suggesting that most actuators function independently with minimal mutual influence.

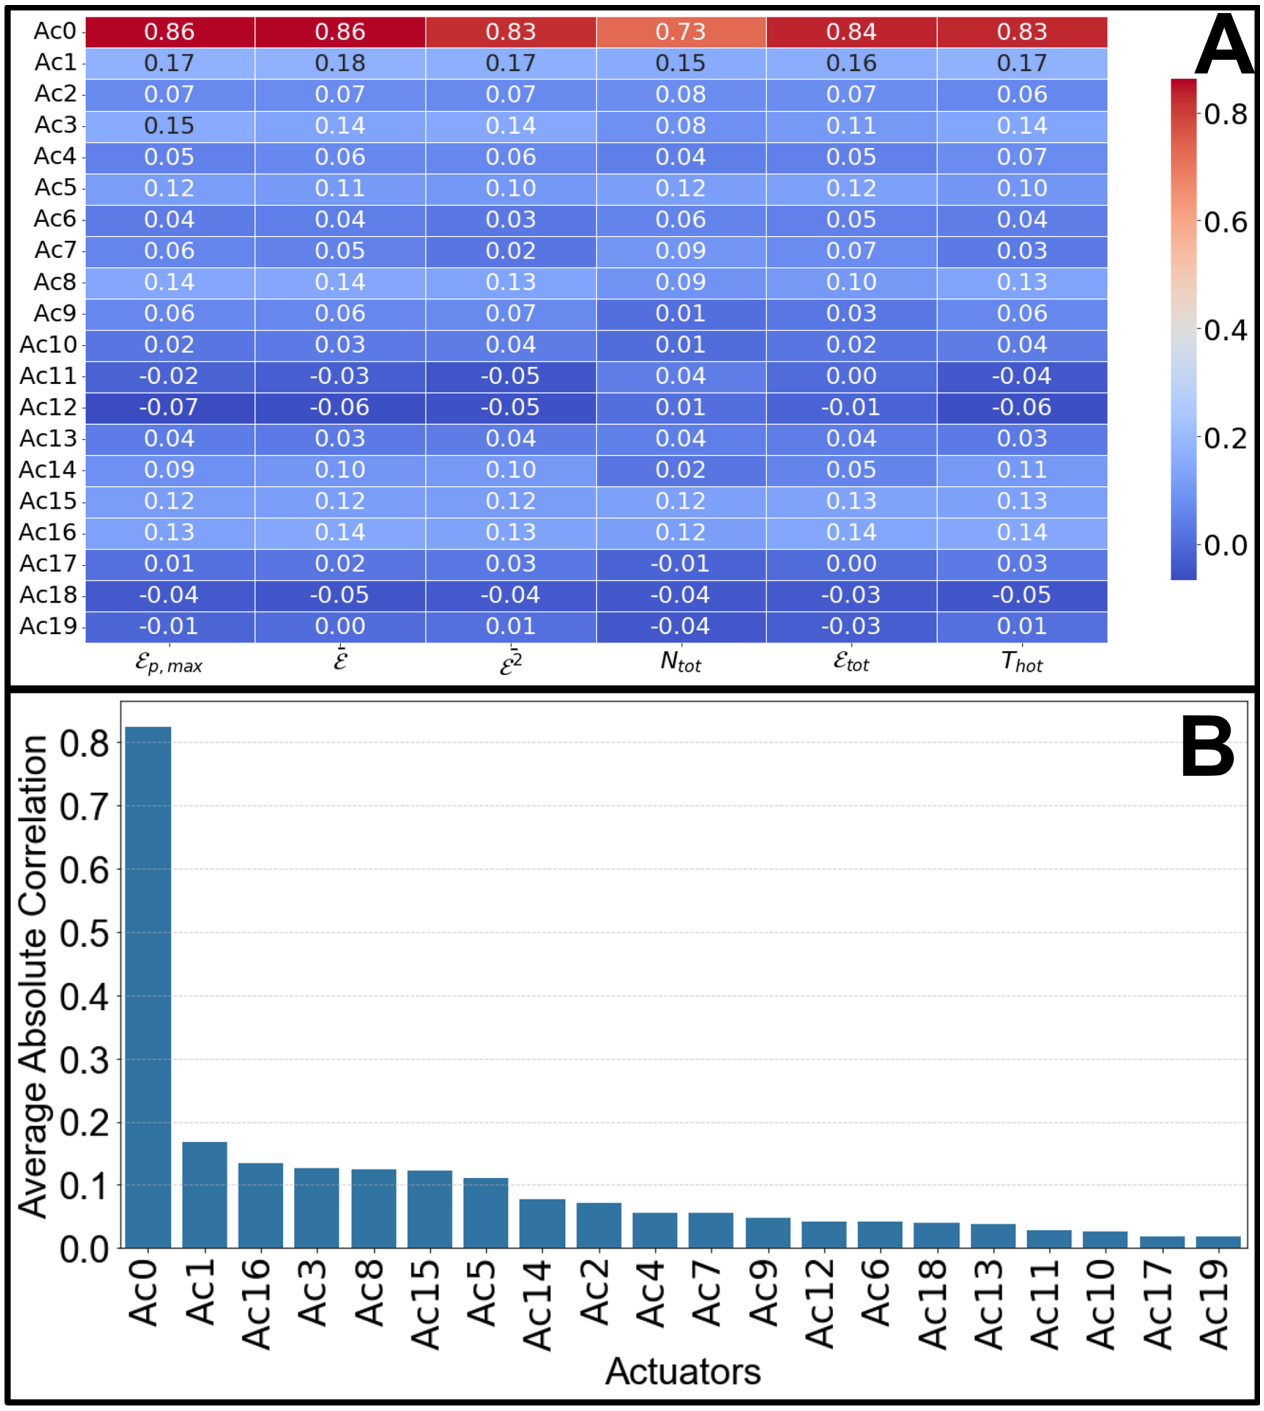

FIGURE 4.10 : Actuator–proton metric correlation analysis.

(A) Correlation matrix heatmap showing the Pearson correlation coefficients between the first 20 actuators of the DM — corresponding to the central region of the mirror surface as defined in Figure 4.7 — and the main proton beam metrics. The values represent the strength and direction of the linear correlation. (B) Average absolute correlation of each DM actuator with all proton beam metrics.

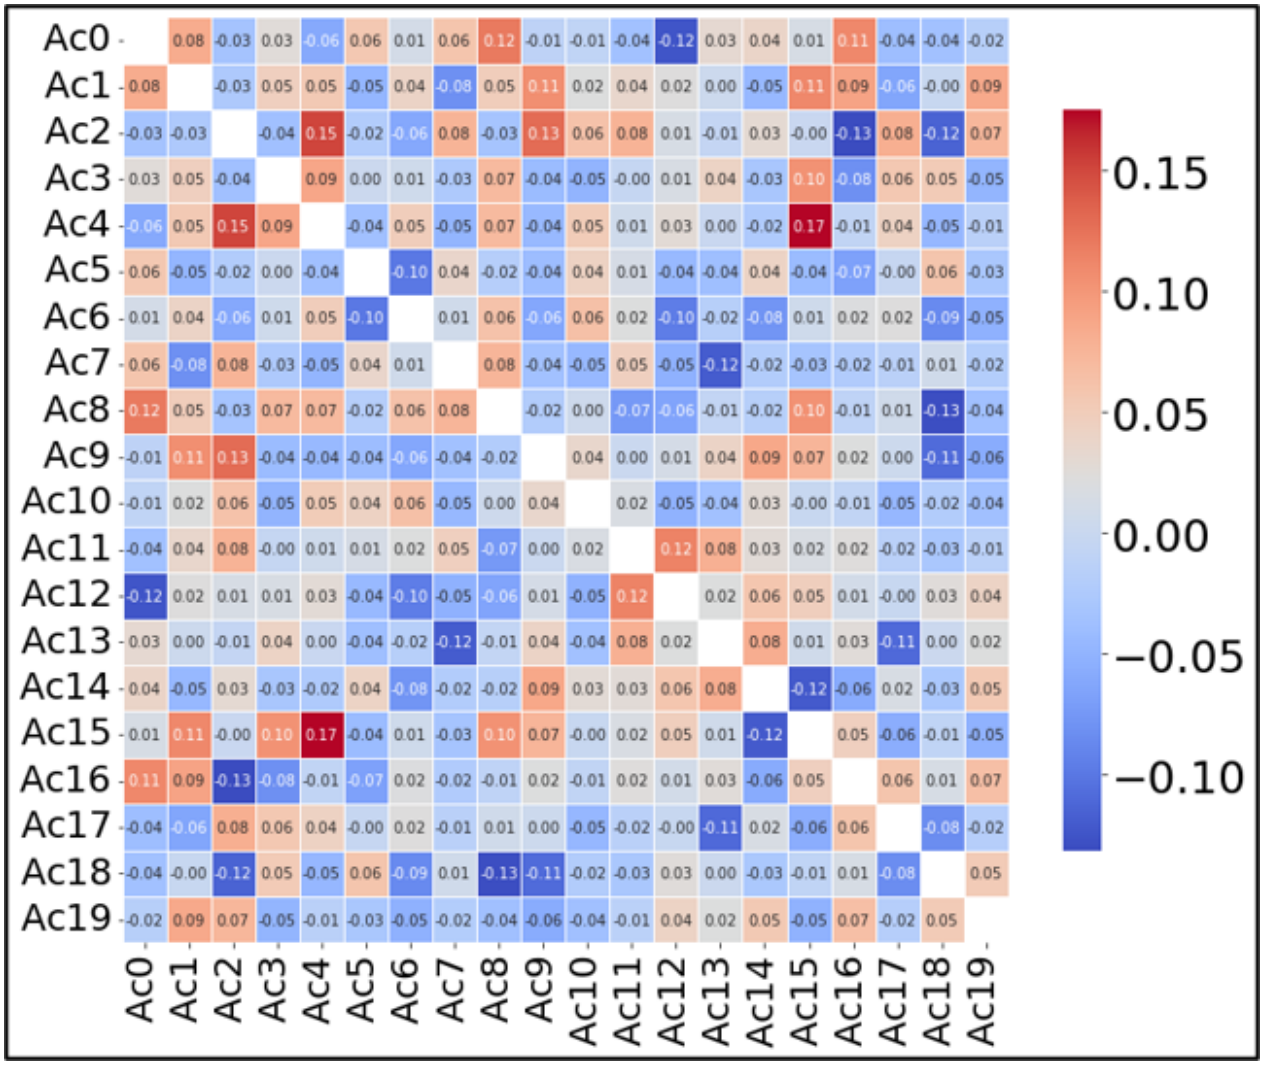

FIGURE 4.11 : Correlation matrix of DM actuators.

Correlation matrix heatmap showing the Pearson correlation coefficients between the 20 actuators of the DM.

Figure 4.12.A presents a correlation analysis among key laser-driven proton beam metrics, showing that evaluated proton beam metrics are nearly all strongly correlated (coefficients $\geq$0.95), indicating that these observables evolve coherently under varying laser beam wavefront conditions. Only N_{tot} exhibits weaker correlations with other metrics (0.73), suggesting a more independent behavior. Figure 4.12.B quantifies the average absolute correlation of each metric with the others, confirming that $\mathcal{E}_{p,max}$ captures the collective behavior of the acceleration process. Unlike integrated metrics, which are more sensitive to shot-to-shot fluctuations and calibration uncertainties, $\mathcal{E}_{p,max}$ exhibits a low standard deviation of only $\sim$5%. This low noise sensitivity, high reproducibility and strong correlation with other proton beam metrics make it a robust indicator of TNSA performance. Its close physical connection to hot electron temperature and therefore sheath field strength further reinforces this aspect (Fuchs et al., 2006). Taken together, these properties make $\mathcal{E}_{p,max}$ a highly informative, stable, and physically meaningful quantity, ideally suited to guide further experimental optimization of laser-driven ion acceleration.

TABLE 4.2 : Sensitivity of proton metrics to wavefront shaping.

Proton beam metrics	$\sigma_{\mathbf{sts}}$	$\sigma_{\mathbf{config}}$	$\sigma_{\mathbf{config}}/\sigma_{\mathbf{sts}}$
N_{tot} [protons·sr^{-1}]	6.85×10^{9}	1.67×10^{10}	2.4
$\mathcal{E}_{tot}$ [J]	2.89×10^{-3}	6.38×10^{-3}	2.2
$\mathcal{E}_{p,max}$ [MeV]	0.37	1.42	3.8
$\bar{\mathcal{E}}$ [MeV]	0.1	0.37	3.7
$\bar{\mathcal{E}^2}$ [MeV^2]	0.7	1.7	2.5

Statistics on the sensitivity of proton beam metrics to laser wavefront shaping. Comparison of shot-to-shot standard deviation ($\sigma_{\mathbf{sts}}$), across-configuration variability ($\sigma_{\mathbf{config}}$), and their ratio for each metric, highlighting the impact of spatial wavefront control on beam performance.

The results presented in Table 4.2 confirm the influence of laser wavefront shaping on key proton beam metrics in the TNSA regime. For all observables, the standard deviation across configurations (σ_{config}) exceeds the shot-to-shot variability (σ_{sts}), with ratios ranging from 2.2 to 3.8. This indicates that spatial modulation of the laser wavefront using the DM introduces systematic and reproducible variations in the beam properties that surpass inherent shot-to-shot fluctuations. Notably, the maximum proton energy ($\mathcal{E}_{p,max}$) exhibits the highest sensitivity, with a ratio of 3.8, underscoring its role as a particularly responsive and informative metric for diagnosing the effects of wavefront control. These findings confirm that spatial shaping of the laser beam is a powerful lever for tuning the coupling conditions at the target surface, directly impacting the efficiency of the acceleration process. Therefore, optimizing the laser wavefront provides a robust and deterministic strategy for enhancing TNSA performance, beyond the limits set by intrinsic laser fluctuations.

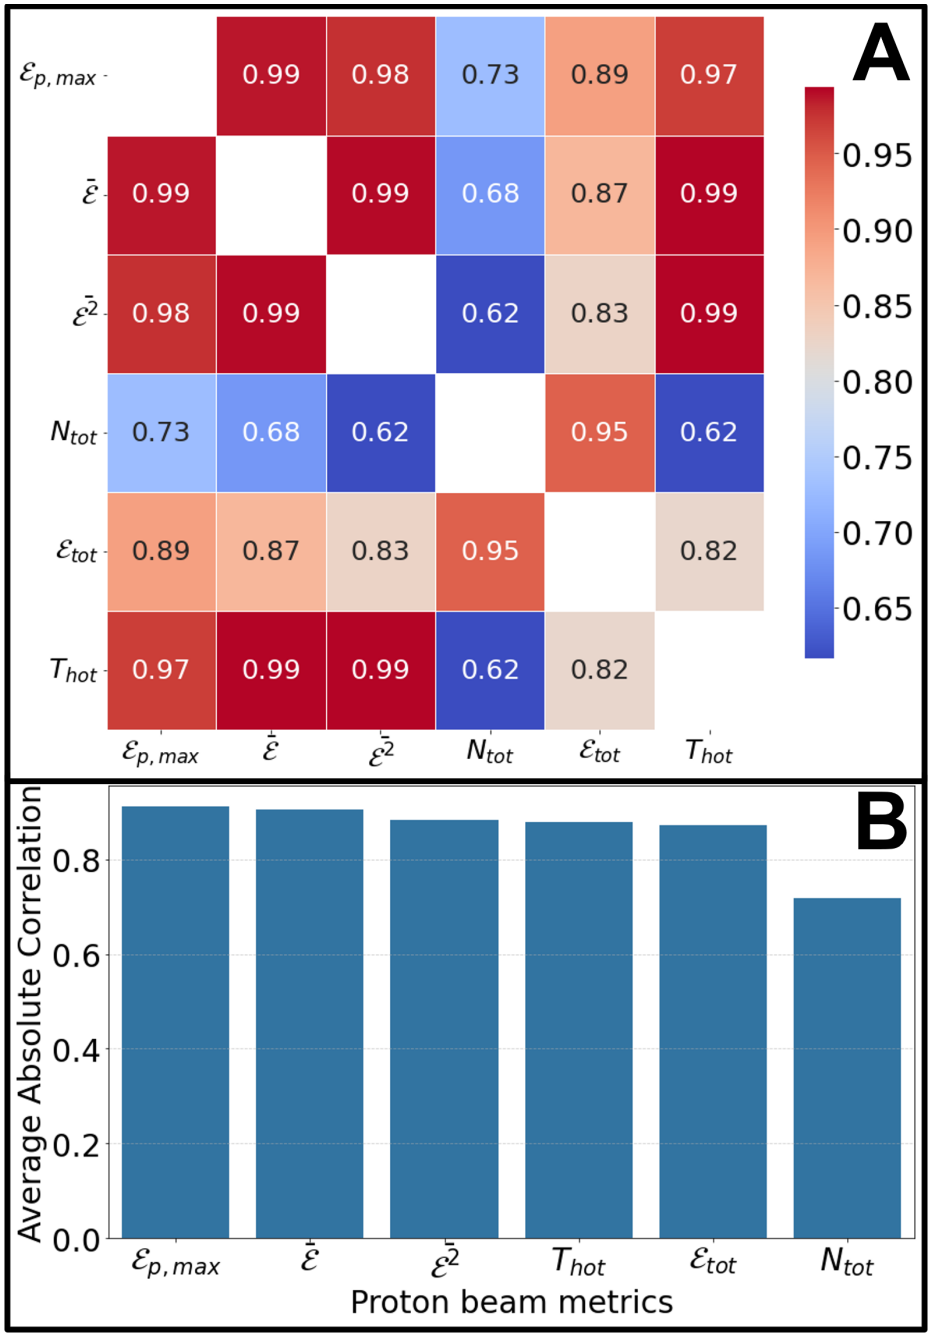

FIGURE 4.12 : Correlation among proton beam metrics.

(A) Correlation matrix heatmap showing the Pearson correlation coefficients among key proton beam metrics measured in the experiment. (B) Average absolute correlation of each proton beam metric with all other beam metrics.

We therefore evaluate the influence of Actuator0 on the maximum proton energy. Figure 4.13 shows the dependence of $\mathcal{E}_{p,max}$ on the voltage applied to Actuator0 of the DM. Each blue point corresponds to a single laser shot, with vertical error bars representing shot-to-shot fluctuations. The red curve shows a second-order polynomial fit to the data. V_0 is the voltage applied to Actuator0. The shaded area around the curve denotes the root mean square error (RMSE) of the fit. The observed nonlinear dependence confirms the important influence of Actuator0 on the acceleration process. This effect is likely mediated through its strong impact on the spatial phase of the laser wavefront, which in turn modulates the intensity distribution at focus and the efficiency of TNSA.

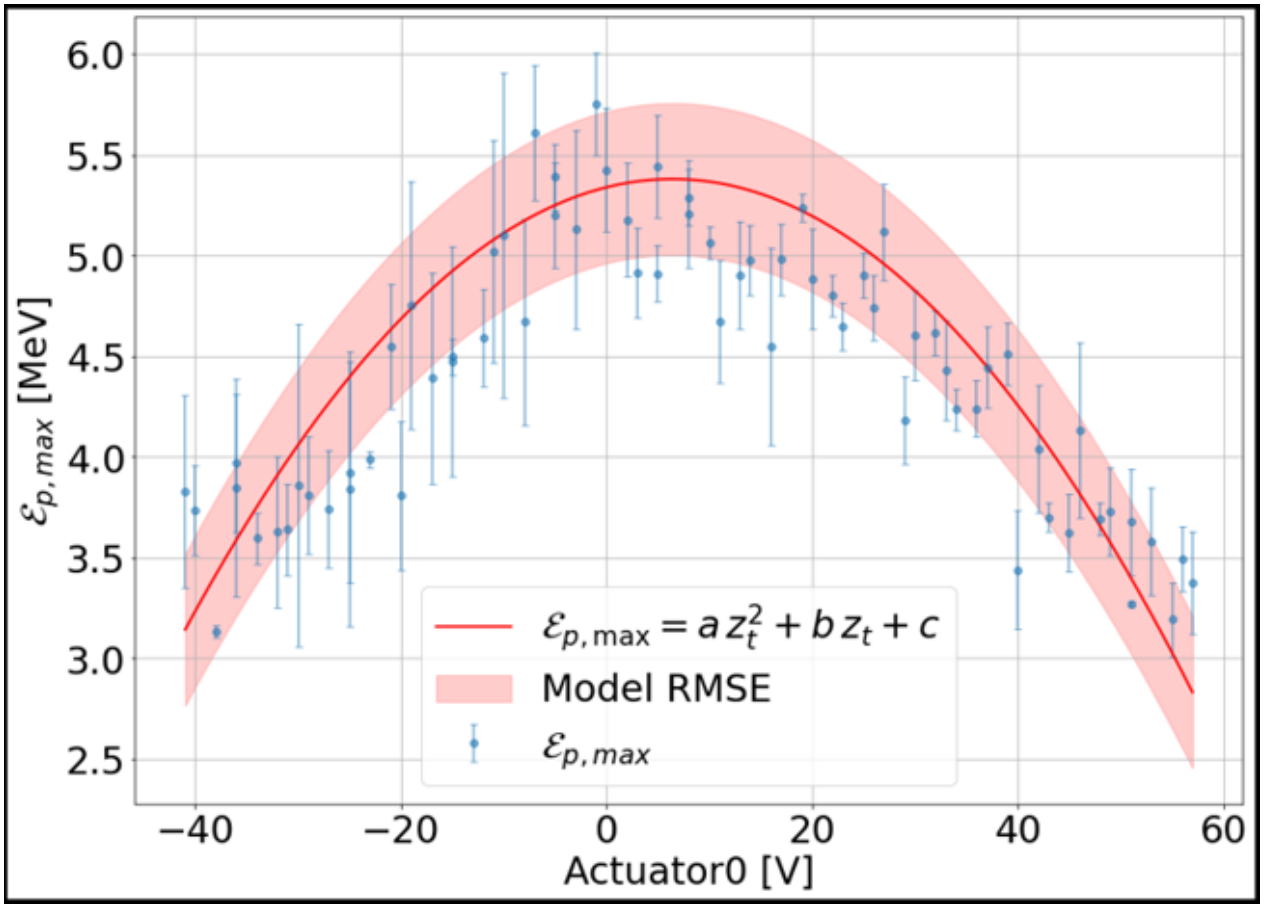

FIGURE 4.13 : Effect of Actuator0 on maximum proton energy.

Each data point corresponds to a single laser shot, with error bars indicating shot-to-shot fluctuations. The red curve is a second-order polynomial fit capturing the nonlinear trend where $a = -9.96 \times 10^{-4}$ MeV.V^{-2}, $b = 1.28 \times 10^{-2}$ MeV.V^{-1}, and $c = 5.34$ MeV. The shaded area represents the model's RMSE.

To capture and interpret the coupled influence of the laser pulse duration (τ_L) and Actuator0 on key proton beam characteristics, we constructed a series of surrogate models using Random Forest (RF) regression. For each proton beam metric, a dedicated RF model was trained on experimental data, using τ_L and the voltage applied to Actuator0 as input features. Each model was configured with 300 decision trees and a fixed random seed to ensure reproducibility. The choice of RF

was motivated by its ability to capture non-linear relationships and handle noisy experimental datasets, which is a known challenge in laser–plasma acceleration experiments. Their robustness to noise and interpretability make them particularly well-suited for data-driven optimization in high-dimensional, low-repetition-rate laser systems. This choice is further supported by the findings of Catrix et al. (2025), where RF models achieved predictive errors close to the intrinsic experimental noise (RMSE $\sim$0.3 MeV), demonstrating their practical value for modeling under realistic conditions in laser-plasma experiments. In the present study, the trained RF models were used to generate two-dimensional contour maps over a densely sampled parameter grid (Figure 4.14.A–E), offering detailed insight into the coupled impact of τ_L and Actuator0 on the proton beam metrics. The models exhibited excellent performance on the training data, with coefficients of determination (R^2) ranging from 0.91 for N_{tot} to 0.94 for $\bar{\mathcal{E}}^2$, and RMSEs that remained low relative to the magnitude of each metric. Across all metrics, a consistent trend emerges : optimal beam performance is achieved in regions where both optimal temporal compression and moderate positive actuator voltages (achieving minimal defocus, i.e., best focused spatially, with a near-planar wavefront at the target plane) are combined. The correlation matrix (Figure 4.14.F) confirms the trends observed in the previous two subsections : the voltage of Actuator0 shows a strong positive correlation with all beam metrics (e.g., $R > 0.6$), whereas the laser's pulse duration exhibits weak or slightly negative correlations.

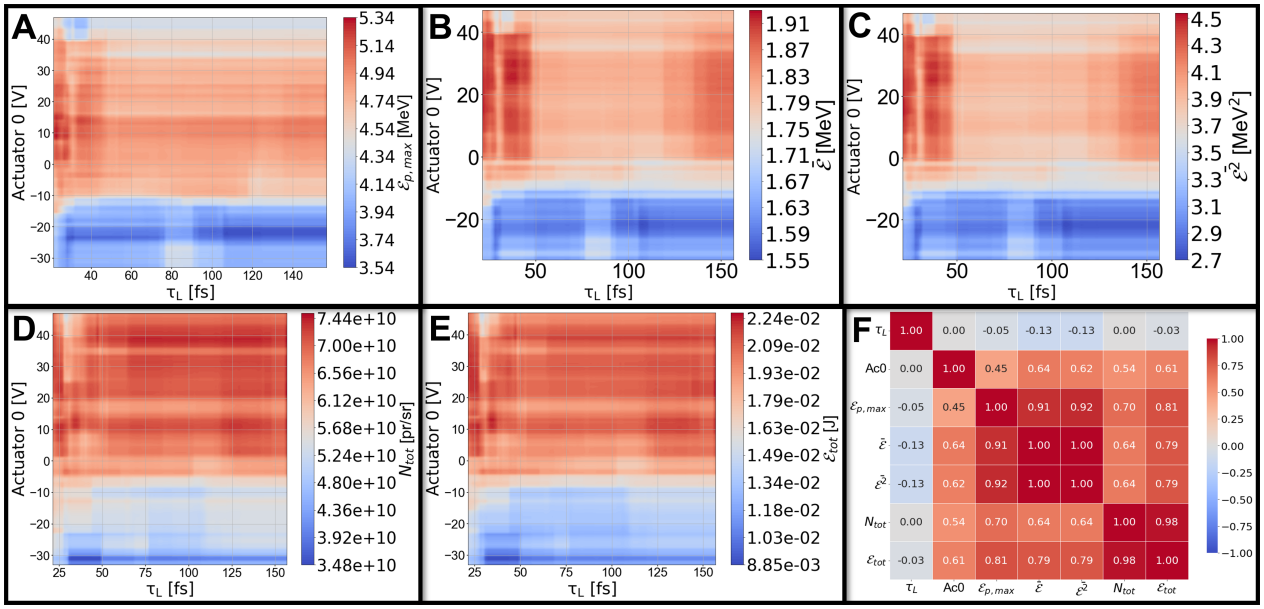

FIGURE 4.14 : Joint influence of τ_L and Actuator0 on proton beams.

(A-E) Two-dimensional maps showing the dependence of proton beam metrics on laser pulse duration τ_L [fs] and the voltage applied to Actuator0 [V] : (A) $\mathcal{E}_{p,max}$, (B) $\bar{\mathcal{E}}$, (C) $\bar{\mathcal{E}}^2$, (D) N_{tot}, and (E) $\mathcal{E}_{tot}$. (F) Correlation matrix between τ_L, Actuator0 voltage, and the proton beam metrics. Strong positive correlations between actuator voltage and all beam metrics are observed, while pulse duration exhibits weak or negative correlations. These findings underscore the dominant influence of spatial wavefront shaping in enhancing TNSA performance.

The analysis reveals that while a subset of actuators, notably Actuator0, exerts a strong influence on both proton yield and energy distribution, pulse duration exhibits only weak or negligible correlations with the proton beam metrics. This suggests that global wavefront control, achieved through coordinated actuator actuation, plays a dominant role in hot electron generation and potentially subsequent sheath field formation in the TNSA regime. As such, optimization strategies should emphasize modal or pattern-based DM control to tailor the focal spot intensity profile for enhanced energy coupling and improved proton beam metrics.

4.7 Discussion

Our comprehensive investigation into spatio-temporal laser pulse shaping reveals a distinct asymmetry in the sensitivity of proton beam metrics to spectral phase versus spatial wavefront control. Adjusting the spectral phase via GDD and TOD showed limited influence on key beam characteristics. As summarized in Table 4.1, the variability in performance metrics across different dispersion settings (σ_{config}) remained comparable to the system's inherent shot-to-shot fluctuations (σ_{sts}). This indicates that, beyond achieving optimal temporal compression, further spectral phase tuning yields marginal benefits under our experimental conditions. In other terms, when varying the pulse duration from 22 fs to $\sim$160 fs while keeping the focal spot and laser energy constant, evaluated proton beam metrics remained approximately constant. Since laser intensity scales inversely with pulse duration ($I \propto 1/\tau_L$), this observation implies a drop in peak intensity over the scan. Figure 4.A of Fuchs et al. (2006) suggests that when pulse duration is increased we slide horizontally to the right across iso-intensity contours (from $\sim 10^{20}$ W.cm^{-2} to $\sim 10^{19}$ W.cm^{-2} in our case), confirming intensity is dropping while $\mathcal{E}_{p,max}$ remains approximately constant. Moreover, in our experiment using a 4.5-μm aluminum foil and a 22 fs, high-intensity laser pulse, hot electrons with an average energy of 1.2 MeV reach the rear target surface within $\sim$16 fs. However, the laser pulse ends before the earliest refluxed electrons can return ($\sim$32 fs), preventing reinforcement of the rear-side sheath field. This is consistent with the model proposed by Mackinnon et al. (2002), which shows that significant enhancement of maximum proton energy only occurs when the laser pulse exceeds the double transit time, enabling sheath densification through electron recirculation. However, no enhancement is observed in our case because the pulse duration is modified by adjusting the GDD and TOD. This approach not only reduces the peak intensity for a fixed pulse energy, but can also degrade the temporal contrast. As illustrated in Osvay et al. (2005), changes to the phase compensation terms — specifically GDD and TOD — can significantly impact the temporal contrast of high-intensity femtosecond pulses. This can introduce preplasma formation or pre-ionization, especially at the target rear, disrupting optimal TNSA conditions. As a result, the expected benefits of longer laser–target interaction time (e.g., enhanced sheath field via reflux or extended acceleration) may be offset by reduced hot-electron temperature, lowered absorption efficiency, and premature target expansion, leading to an observed plateau in maximum proton

energy as observed. This suggests that, within the explored range, proton acceleration efficiency is not solely determined by laser intensity.

In contrast, spatial wavefront shaping via the DM demonstrated a clear and consistent influence on all major beam metrics. Among the 20 actuators studied, Actuator0 emerged as the dominant control parameter, showing strong correlations with $\mathcal{E}_{p,max}$, $\bar{\mathcal{E}}$, $\mathcal{E}_{tot}$, and T_{hot}. It also exhibited the highest absolute average correlation with all other proton beam metrics, confirming its pivotal role in modulating laser–plasma coupling. Our study also suggests that the system is primarily sensitive to a reduced subset of actuators, allowing efficient optimization within a lower-dimensional parameter space. Identifying the most influential actuators is therefore critical for reducing experimental and computational complexity.

$\mathcal{E}_{p,max}$ shows both high correlation with all other proton beam metrics and low relative shot-to-shot fluctuation ($\sim$5%), making it a reliable and informative optimization target. Therefore, it can serve as a practical proxy for enhancing the overall efficiency of the TNSA process. Since $\mathcal{E}_{p,max}$ reflects the strength of the sheath field and the efficiency of hot electron generation, optimizing this parameter inherently improves the collective acceleration dynamics.

The apparent difference in the correlation of Actuator0 with the proton beam metrics between Figure 4.10.A and Figure 4.14.F arises from how Pearson's correlation coefficient captures statistical variance under different experimental designs. Pearson's correlation $\rho_{X,Y}$ is defined as :

Equation 4.20 : Pearson correlation coefficient between two variables X and Y

$$\rho_{X,Y} = \frac{\sum_i (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_i (X_i - \bar{X})^2} \sqrt{\sum_i (Y_i - \bar{Y})^2}} \tag{4.20}$$

and measures the degree of linear co-variation between two variables. In the multi-actuator scans (Figure 4.10.A), where all actuators vary independently following a LHS, the influence of Actuator0 on the proton beam metrics is largely monotonic. This leads to an elongated elliptical distribution in the (V_0, Y) plane (where Y is a given proton beam metric), producing consistently high values of $(X_i - \bar{X})(Y_i - \bar{Y})$ in the numerator of $\rho_{X,Y}$, and thus a strong linear correlation. In contrast, in the 2D scans (Figure 4.14.F), where τ_L is varied simultaneously with Actuator0, the response becomes more non-monotonic and symmetric, often peaking around $V_0 \approx 3$ V and decreasing on both sides. This introduces two sign changes in $(Y_i - \bar{Y})$ across the range of V_0, causing positive and negative terms in the covariance sum to cancel out, thereby reducing the correlation even if Actuator0 still drives the main physical variation. As a result, Pearson's coefficient underestimates the true influence of Actuator0 in designs where the response is nonlinear.

These findings highlight the importance of rethinking optimization strategies for laser-driven proton acceleration. Although laser peak intensity is commonly regarded as the main factor governing hot

electron temperature and sheath field strength, our results — together with prior studies (Fourmaux et al., 2011b; Esirkepov et al., 2014; Yao et al., 2024b) — indicate that excessive intensity can be detrimental. Temporal artifacts such as pre-pulses or imperfect compression may generate underdense pre-plasma, reducing energy coupling efficiency. Moreover, nonlinear propagation effects, including relativistic self-focusing (Sprangle et al., 1987), filamentation (Metternich et al., 2022), and reflectivity induced by plasma instabilities (Dover et al., 2020a; Gizzi et al., 2021) can disrupt energy confinement and broaden the hot electron distribution. As a result, a moderately intense, well-compressed, and high-contrast pulse may outperform higher-intensity configurations in terms of reproducibility and acceleration efficiency (Takagi et al., 2021; Willim et al., 2023).

Altogether, these results highlight the necessity of systematic parametric studies and data-driven strategies to navigate the complex, nonlinear landscape of TNSA. Achieving optimal performance requires more than maximizing intensity; it demands fine-tuned control over the laser's spatio-temporal profile to ensure stable, efficient coupling to the plasma and the generation of high-quality ion beams.

4.8 Conclusion

Our results demonstrate that spatial wavefront shaping, particularly through targeted actuator control, has a significantly greater impact on proton beam quality than spectral phase adjustments around optimal compression under our experimental conditions. Among all metrics, the maximum proton energy ($\mathcal{E}_{p,max}$) emerged as a sensitive and reliable indicator of acceleration performance, serving as a robust optimization target. These findings underscore the importance of spatio-temporal laser shaping and support a shift from intensity-driven strategies toward more refined control of laser–plasma coupling conditions to enhance TNSA efficiency.

5 20-DIMENSIONAL SURROGATE-ASSISTED BAYESIAN OPTIMIZATION OF LASER-DRIVEN PROTON BEAMS

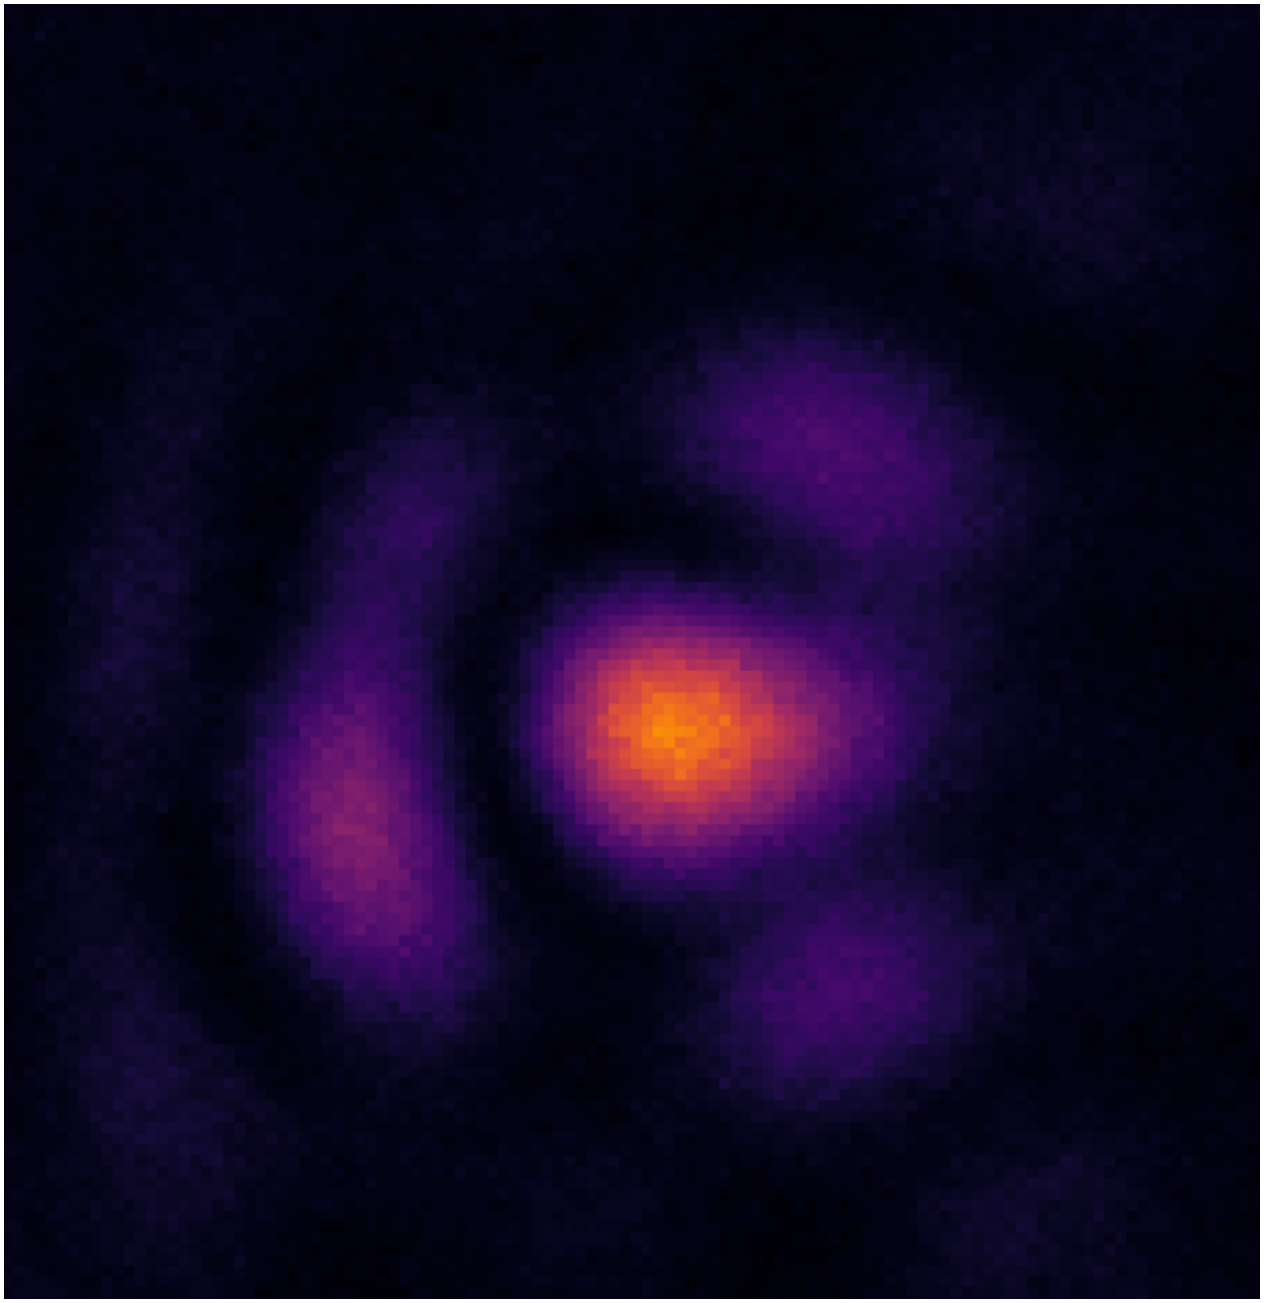

Spatial intensity profile of the laser focal spot. Image originale réalisée par l'auteur (Elias Catrix). Original image created by the author (Elias Catrix). No external permission is required; use authorized by the author.

Foreword — This chapter distills and extends the results recently reported in *Applied Physics Letters* named "20-dimensional surrogate-assisted Bayesian optimization of laser-driven proton beams." (Catrix et al., 2025).

In this chapter, by leveraging Machine Learning (ML) to optimize the laser wavefront, we demonstrate the improvement in key performance metrics, including maximum proton energy ($\mathcal{E}_{p,max}$), mean proton energy ($\bar{\mathcal{E}}$), total proton beam energy ($\mathcal{E}_{tot}$), integrated proton yield (N_{tot}), laser-to-proton energy conversion efficiency (α), second order moment (M^2) and hot electron temperature (T_{hot}). This was accomplished through adaptive control of a deformable mirror (DM) using a multi-step Random Forest surrogate-assisted Bayesian optimization approach. Starting from zeroed DM actuator voltages, our method identified an optimal configuration using 20 out of 48 actuators, requiring fewer than 150 experimental data samples. Our method outperformed conventional wavefront correction, which typically minimizes aberrations to converge toward a flat wavefront by leveraging real-time feedback from a wavefront sensor, with the primary aim of increasing the proton beam energy. Integrating advanced wavefront control, this data-driven method challenges the preference for correcting aberrations to achieve a flatter wavefront in laser-driven ion acceleration. We also propose a strategy for optimizing short focal length laser-driven ion accelerators at facilities where measuring the wavefront at nominal full laser power is not implemented.

5.1 Introduction

Among laser-driven ion acceleration techniques, Target Normal Sheath Acceleration (TNSA) permits the production of ultra-short proton bursts with high peak energies (Wilks et al., 2001). However, achieving stable and repeatable high-brightness beams remains challenging due to an unstable laser-plasma interaction that impacts particle yields, energies, and the divergence profile. Precise control of laser parameters is essential to optimize these complex interactions.

Machine Learning (ML) has recently gained interest as a data-driven approach for exploring complex parameter spaces in experimental physics (Anirudh et al., 2023; Döpp et al., 2023; Roussel et al., 2024). By efficiently guiding the search for optimal conditions, it offers a promising avenue for improving laser-based particle acceleration. This approach leverages real-time diagnostics and predictive modeling to refine experimental parameters with minimal iterations, i.e, in a minimal number of laser shots. Recent developments have highlighted the growing role of ML in modeling and optimizing laser–plasma interactions, including neural network-based surrogate modeling for phase-stable acceleration dynamics (Liu et al., 2024) and synthetic diagnostics predicting on-shot proton energy spectra using deep variational autoencoders and ensemble learning strategies (McQueen et al., 2025).

In particular, Bayesian optimization (BO) (Jones et al., 1998; Mockus, 2005; Brochu et al., 2010; Shahriari et al., 2015) has proven effective for optimizing computationally or experimentally intensive, noisy, black-box objective functions in low-dimensional parameter spaces (typically ≤ 20)

(Frazier, 2018). BO builds a probabilistic model, usually a Gaussian Process (GP), of the objective function - i.e, the function to optimize - allowing efficient sampling of points. This probabilistic BO technique has been used to maximize the pulse energy at the LCLS x-ray free-electron laser facility (Duris et al., 2020) and improve the electron beam quality in laser wakefield electron acceleration (Shalloo et al., 2020; Jalas et al., 2021, 2023). However, its application for laser-driven proton acceleration remains relatively underexplored (Loughran et al., 2023), largely due to the limited repetition rates (Nishiuchi et al., 2009; Gao et al., 2017; Treffert et al., 2022a; Giuffrida et al., 2023; Nedbailo et al., 2023; Catrix et al., 2023) of ion beamlines and the noisy data resulting from non-linear laser-plasma interactions (Sprangle et al., 1990; Kaw, 2017).

The laser spatial spot size and shape are crucial for efficient laser-driven ion acceleration and depend on the wavefront profile. Conventional wavefront correction methods, guided by wavefront sensors, can improve beam quality (Beaurepaire et al., 2015; Lin et al., 2019; Kurz et al., 2021), often optimizing towards Gaussian profiles and flat wavefronts. However, several groups have demonstrated superior performances using non-Gaussian beams (Brabetz et al., 2015; Willim et al., 2023; Wang, 2024), hence, these wavefront correction methods face challenges in navigating the complex parameter space of ion acceleration as they don't directly optimize the final ion beam parameters.

In this chapter, we present a data-driven framework that integrates advanced wavefront control with a multi-step BO (MSBO) strategy to enhance the overall performance of laser-driven proton beams in the TNSA regime. By adaptively adjusting the actuators of a DM, this approach systematically explores non-intuitive wavefront profiles, challenging traditional assumptions about optimal laser beam profiles for TNSA ion acceleration. Through experimental feedback, this method improves multiple beam characteristics. Moreover, the technique is well-suited for high-intensity laser systems that lack terawatt or petawatt attenuators, where direct focal spot characterization at full power is infeasible. It offers a practical, model-free solution for optimizing interaction conditions under realistic operational constraints.

5.2 Experimental setup

Experiments were conducted on the laser-driven ion acceleration beamline at the *Institut national de la recherche scientifique* (*INRS*) in Varennes, Canada, using the *Advanced Laser Light Source (ALLS)* 150 TW Ti : Sapphire laser (Vallières et al., 2020). This double chirped-pulsed amplified (CPA) system delivers 3.2 J on target with a pulse duration of $\tau_{\mathrm{FWHM}} = 22$ fs (Full-Width-Half-Maximum) and a central wavelength of $\lambda_0 = 800$ nm. The experimental setup, illustrated in Figure 5.1, employs an $f/3$ off-axis parabola (OAP) to focus the 95 mm diameter beam (measured at e^{-2}) down to a focal spot size of $w_{\mathrm{FWHM}} \sim 5$ μm, achieving a peak intensity I_0 of approximately 1.3×10^{20} W.cm^{-2}. Before laser pulse compression the beam underwent a cross-wave polarizer (XPW) cleaning, ensuring an Amplified Spontaneous Emission (ASE) pre-pulse contrast $< 10^{-10}$

at -100 ps relative to the main pulse and a steep power rise with a contrast $< 10^{-6}$ at -3 ps (see Figure 5.2). A piezoelectric bimorph DM with 48 actuators from *AKAOptics* (with a maximum voltage range [-200V, +200V]), shown in Figure 5.3, allowing for precise wavefront shaping, and a *SID4* wavefront sensor (WFS) from *Phasics* provided real-time feedback to assess quantitatively the impact of DM actuator configurations on the laser beam wavefront. At the same time, a CCD camera was used for beam visualization at the focus position. This beam analysis is performed at the laser nominal energy of 3.2 J by attenuating the laser beam energy with wedges before the diagnostics.

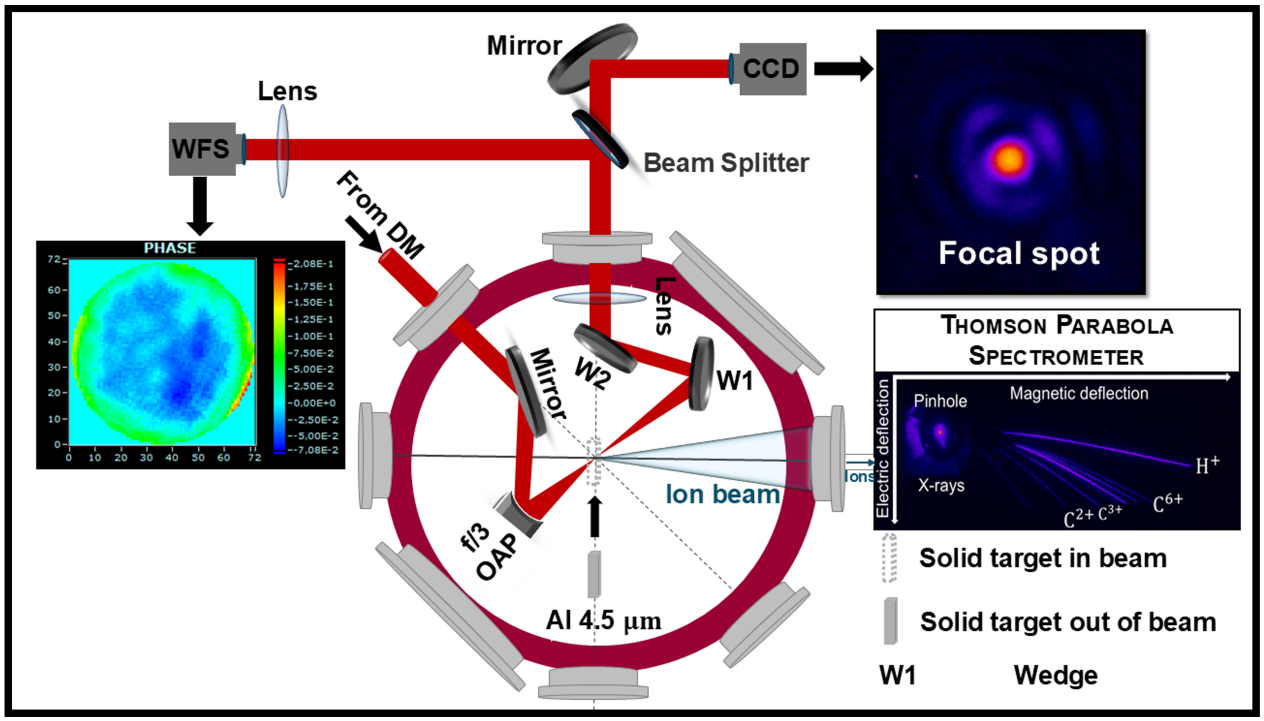

FIGURE 5.1 : Experimental setup layout.

Sketch of the experimental setup including direct visualization of the laser focus (CCD) and laser wavefront real-time diagnostics (WFS), both measured at maximum laser energy (3.2 J). This figure has been adapted from Catrix et al. (2025), with the publisher's permission.

The experiment utilized p-polarized laser pulses incident at a 20° angle relative to the target normal on a 4.5-μm-thick aluminum foil to generate ion beams via the TNSA regime. The target, mounted on a multi-target holder accommodating up to 400 targets, was positioned at the laser focus, enabling systematic and efficient experimentation at a repetition rate of 0.625 Hz (Catrix et al., 2023). A Thomson parabola spectrometer (TPS) coupled with a microchannel plate (MCP) detector was employed to measure the proton spectrum and record the proton beam metrics at every laser shot. More information on the setup can be found in Vallières et al. (2020) and Catrix et al. (2023).

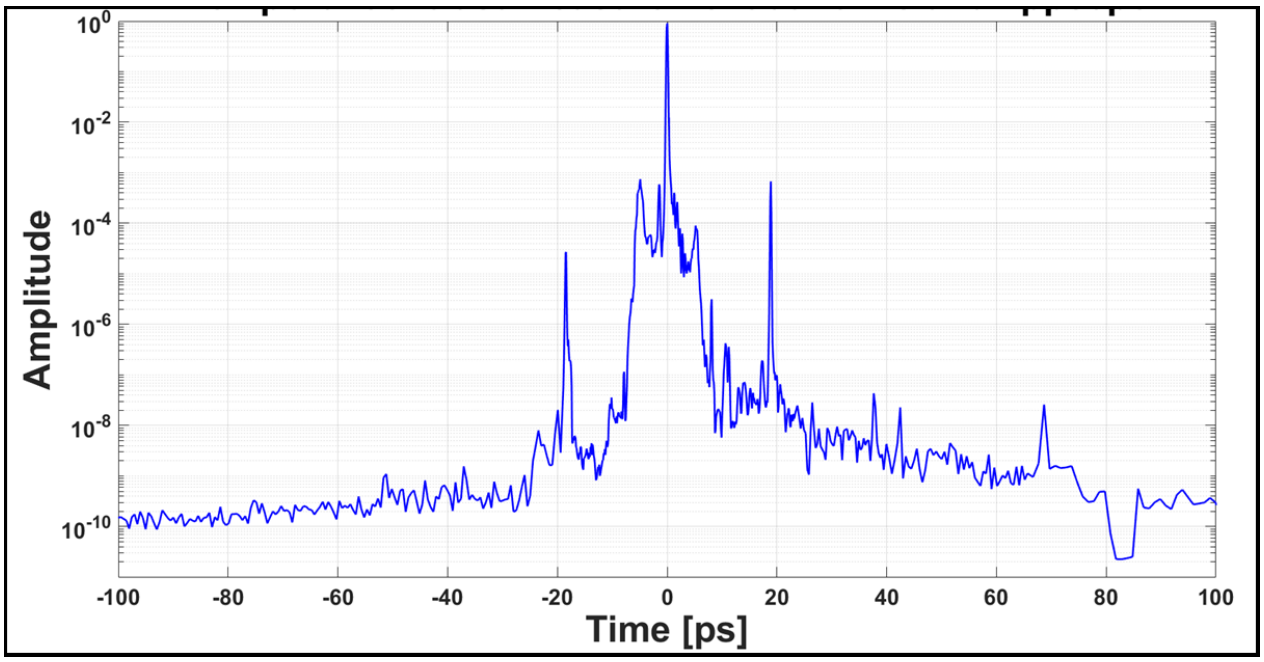

FIGURE 5.2 : Temporal contrast of the *ALLS* 150 TW laser.

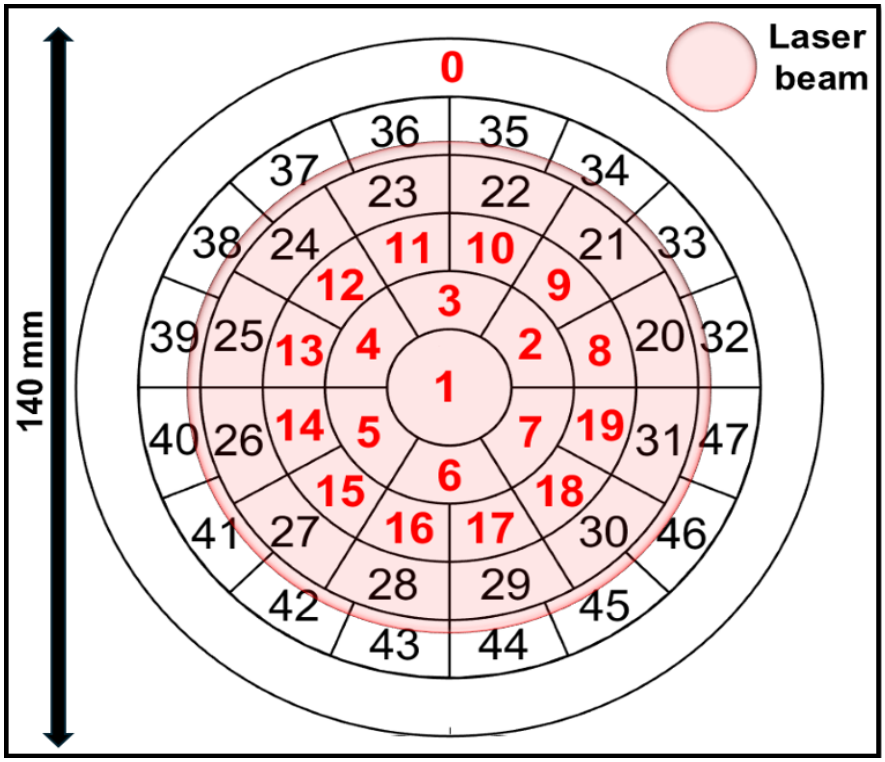

FIGURE 5.3 : Layout of the 48 actuators on the DM.

The 48 actuators on the deformable mirror are used for laser wavefront control. Actuators are numbered from 0 to 47, with those utilized for the analysis highlighted in red (actuators 0–19). Actuator 0 corresponds to the global defocus mode, which translates the entire mirror surface forward or backward.

5.3 Methodology

To optimize the laser-driven proton beams, we aimed to optimize the maximum proton energy, chosen for its broad dynamic range and strong sensitivity to laser wavefront conditions, as detailed in Chapter 4. This beam property provides a clear distinction between tested DM configurations and offers a reliable metric for optimization. To leverage this, we implemented an MSBO strategy that integrates a Random Forest (RF) model with BO, as illustrated in Figure 5.4. RF was selected for its ability to model complex non-linear relationships and perform reliably in the presence of noise and high-dimensional data. Its effectiveness in both predictive accuracy and interpretability has been demonstrated in applications such as disruption classification and hybrid physics–ML frameworks for plasma stability forecasting (Rea et al., 2018; Speiser et al., 2019; Piccione et al., 2020). In our study, the RF model was trained to approximate the relationship between DM actuator voltages and the resulting maximum proton energy. BO was then employed to efficiently explore the actuator parameter space and identify configurations that maximize the predicted energy. This combined approach led to the identification of an optimized DM configuration with performance stable within experimental uncertainty bounds.

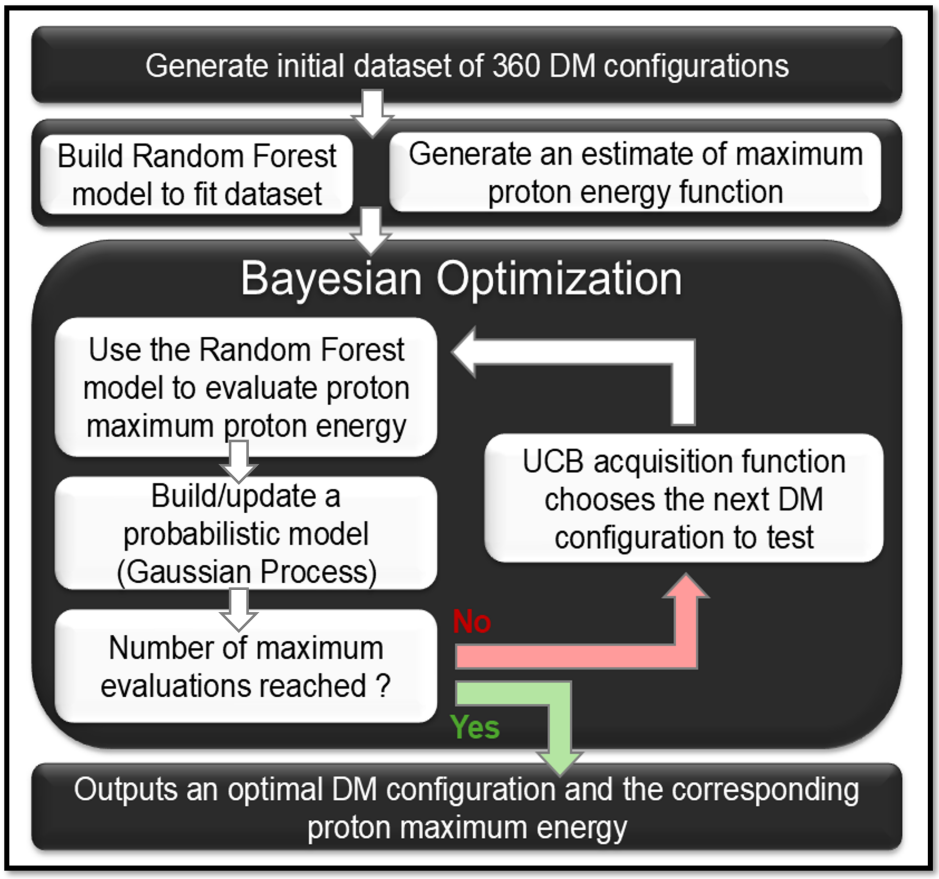

FIGURE 5.4 : Multi-Step Bayesian Optimization workflow for DM tuning.

This figure has been adapted from Catrix et al. (2025), with the publisher's permission.

Step 1

Generate initial dataset of 360 DM configurations

Example of data points:

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	$E_{p,max}$
3	-43	-4	-44	-43	68	6	37	47	-29	39	-13	-55	-49	-27	-77	10	34	14	7	3.64
50	60	-58	54	8	44	-55	-79	-57	-24	-68	2	23	-77	16	71	-69	-75	12	-44	5.17

FIGURE 5.5 : Step 1 of the MSBO approach.

The first step (Figure 5.5) in the optimization process involves generating an initial dataset to train a surrogate model. Starting from a reference configuration, where all 48 DM actuators are set to 0 V and saved as a .txt file compatible with the DM control system, voltage variations were introduced on a selected subset of actuators. Indeed, to reduce dimensionality and focus on the most impactful region, only the central 20 actuators (out of 48) were varied, based on their proximity to the high-intensity area of the laser focal spot. This selection was aimed at improving computational efficiency, as BO is most effective for parameter space dimensions typically $\leq$ 20 (Frazier, 2018). Each of the 20 actuator voltages was allowed to vary within $\pm$80 V relative to the reference configuration (all DM actuators set to 0 V). To ensure uniform and space-filling sampling of the multidimensional parameter space, Latin Hypercube Sampling (LHS) was employed (Helton et al., 2003; Shields et al., 2016). Unlike simple random sampling, LHS guarantees that each actuator samples its range evenly, avoiding clustering or gaps. The *lhsdesign* function in *MATLAB* was used to generate these configurations efficiently. The generated actuator configurations were stored in an *Excel* spreadsheet, with each row representing a unique DM voltage map. These were later exported as .txt files for direct use with the DM control system during experimental runs for evaluation of the proton beam metrics. An illustration of this dataset construction process is shown in Figure 5.5. Finally, a correlation analysis between the actuators and the measured proton beam parameters revealed that Actuator0, responsible for controlling laser beam defocus, had the highest correlation (0.86) with the maximum proton energy, indicating its dominant influence. In contrast, most other actuators exhibited weak correlations (average of 0.08), and inter-actuator correlations remained low ($< |0.1|$), confirming that the actuators operate largely independently.

We then evaluated several surrogate models to identify the most suitable one for capturing the relationship between the DM actuator voltages and the resulting maximum proton energy. Table 5.1 presents a comparative summary of five ML regression models, highlighting their performance in terms of the coefficient of determination (R^2) and the root mean square error (RMSE). The RMSE is a commonly used metric in ML indicating the model's predictive performance. It is defined as :

Equation 5.1 : Root Mean Square Error of a regression model

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{5.1}$$

where n is the number of laser shots, y_i are the actual (experimentally measured) maximum proton energy values, and $\hat{y}_i$ are the predicted maximum proton energy values.

The models span linear, ensemble-based, instance-based, kernel-based, RFs, and neural network paradigms. The RF model emerged as the most suitable surrogate for our application, achieving an R^2 of 0.97 and an RMSE of 0.24 MeV. This indicates a high predictive accuracy and excellent generalization over the observed data without signs of overfitting. RFs are particularly robust for modeling complex nonlinear relationships and handling high-dimensional input spaces, which makes

them well-matched to the 20 actuator parameters involved in our experiment. The Linear Regression and Multi-layer Perceptron (MLP) regressor models also showed reasonable performance, with R^2 values of 0.81 and 0.78, respectively. However, the linear model cannot capture nonlinear effects, while the MLP requires more data to reliably train and may exhibit instability without extensive hyperparameter tuning. In contrast, K-Nearest Neighbors (KNN) and Support Vector Regressor (SVR) underperformed, with lower R^2 values and higher RMSEs.

Among the tested models, KNN and SVR algorithms were found to be more sensitive to feature scaling and local fluctuations in the parameter space. This is primarily because both methods rely on distance- or similarity-based metrics to define relationships between data points. In KNN, variations in the relative magnitude of features or the presence of noise can distort the neighborhood geometry, leading to unstable predictions when the data are not properly normalized. Similarly, in SVR, the kernel function (typically radial or polynomial) depends on pairwise distances, making the regression surface highly sensitive to sharp variations and local irregularities in the input space. In contrast, ensemble and probabilistic models such as RF or NN demonstrate greater robustness, as they learn adaptive internal representations or smoothing scales that mitigate the effects of scaling and noise.

Based on this comparative analysis, we selected the RF model as the surrogate of choice for guiding optimization in our MSBO framework. Its balance between interpretability, robustness, and predictive power makes it well-suited for modeling the complex parameter space of laser-driven ion acceleration experiments.

TABLE 5.1 : Comparison of surrogate models for $\mathcal{E}_{p,max}$ prediction.

Model	**Description**	$\mathbf{R^2}$	**RMSE [MeV]**
Linear Regression	A simple model that assumes a linear relationship between inputs and the target.	0.81	0.62
Random Forest	An ensemble of decision trees that improves prediction by averaging multiple trees.	0.97	0.24
Support Vector Regressor	A kernel-based method that tries to fit data within a specified error margin.	-0.022	1.40
K-Nearest Neighbors	Predicts values based on the average of k-nearest neighbors in the training set.	0.68	0.79
MLP Regressor	A neural network model that uses multiple layers to capture nonlinear relationships.	0.78	0.67

Therefore, in the second step of our MSBO (Figure 5.6), an RF model is trained on the initial dataset to estimate the relationship between the 20 DM actuators and the proton beam metrics. Using 300 decision trees, predictions are made with robustness and noise reduction. In the algorithm, the trees grow until the leaves are pure. A pure leaf, in this context, means that a set of actuator voltages results in the same proton energy output. Figure 5.7 compares the actuator importance profiles derived from RF regressors trained independently on six different proton beam metrics :

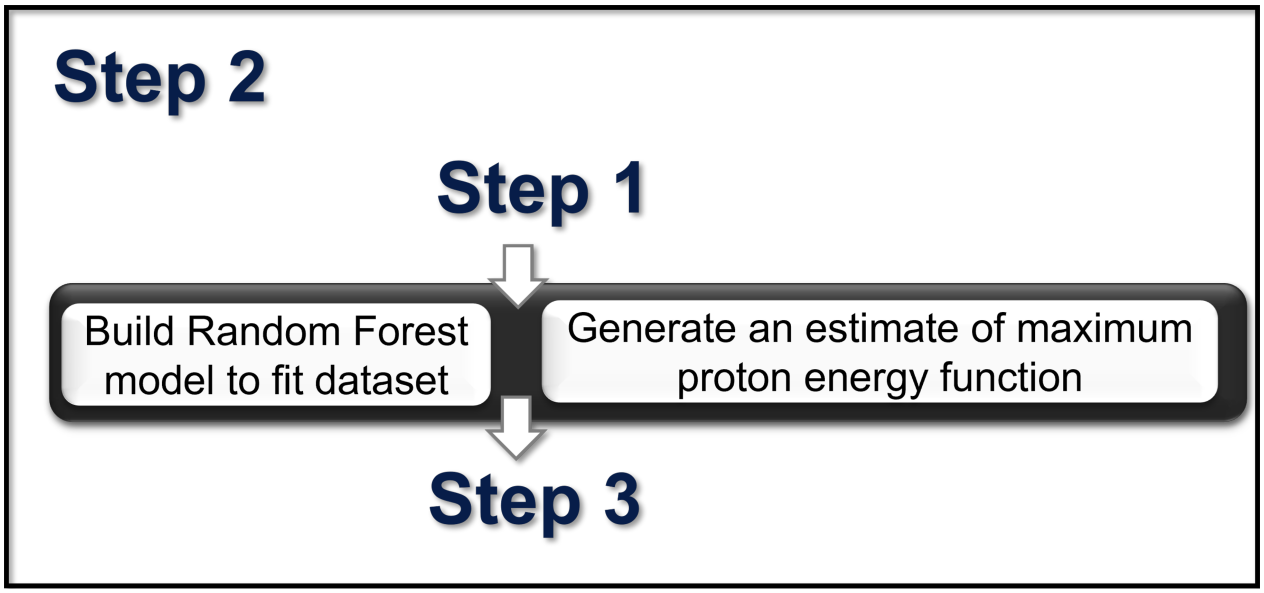

FIGURE 5.6 : Step 2 of the MSBO approach.

maximum proton energy ($\mathcal{E}_{p,max}$), mean proton energy ($\bar{\mathcal{E}}$), total proton beam energy ($\mathcal{E}_{tot}$), integrated proton yield (N_{tot}), second-order moment (M^2) and hot electron temperature (T_{hot}). Each line represents how strongly the 20 DM actuators influence the prediction of a given metric. The curves reveal a consistent pattern across all outputs, indicating that the same subset of actuators (notably Actuator0) plays a dominant role in shaping all beam properties. Importantly, among the evaluated metrics, $\mathcal{E}_{p,max}$ is both equally sensitive to actuator tuning and exhibits the lowest shot-to-shot fluctuations, typically within $\pm$5% RMS, whereas other metrics such as N_{tot} or $\mathcal{E}_{tot}$ fluctuate by $\pm$15–25%. This confirms that the proton maximum energy provides a practical and reliable target for optimization. Constraining the BO process to this metric ensures rapid convergence and meaningful experimental gain, while still indirectly improving correlated metrics due to the shared actuator dependencies observed in this plot and insights given in Chapter 4. Figure 5.8 displays the predictive performance of RF regression models trained to estimate six distinct beam output metrics using the 20 actuator voltages of the DM as input features. Each subplot shows predicted values plotted against measured ground truth, with the ideal 1 :1 correspondence indicated by a dashed red line. Across all proton beam metrics, RF achieves high accuracy, with coefficients of determination $R^2 \geq 0.95$. This strong agreement indicates that the actuator configurations encode sufficient information to accurately predict not only global beam energy but also spectral and thermal properties of the proton source. These results validate the use of RFs as effective surrogate models for downstream optimization.

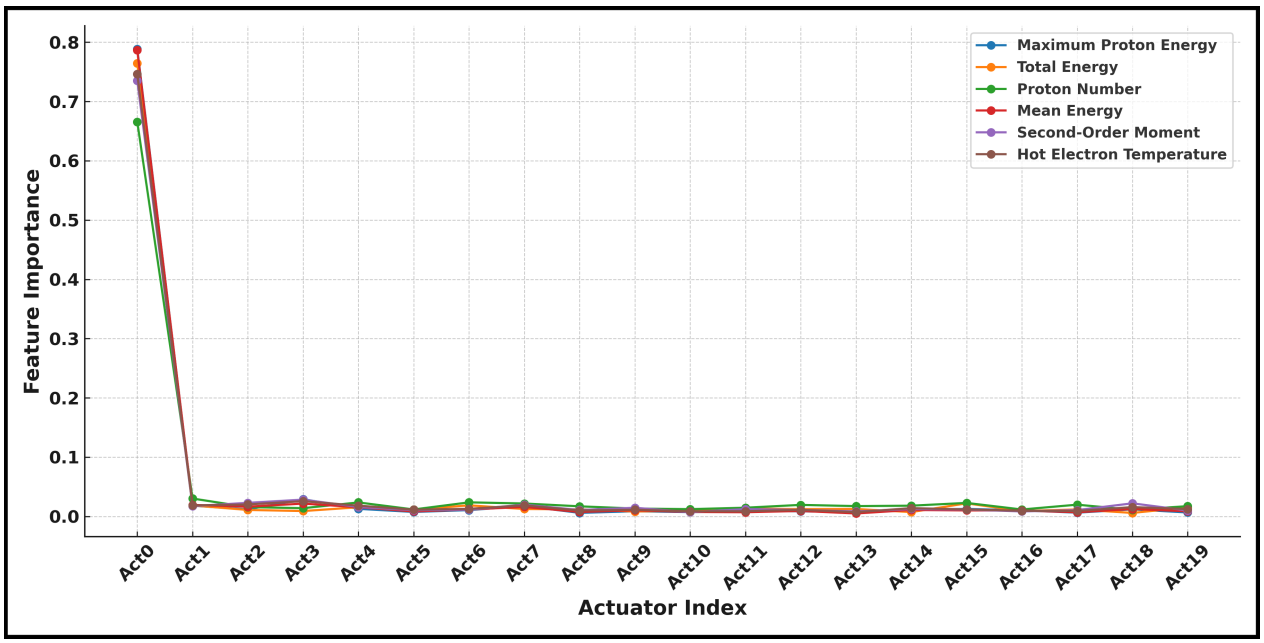

FIGURE 5.7 : Actuator feature importance for proton beam metrics.

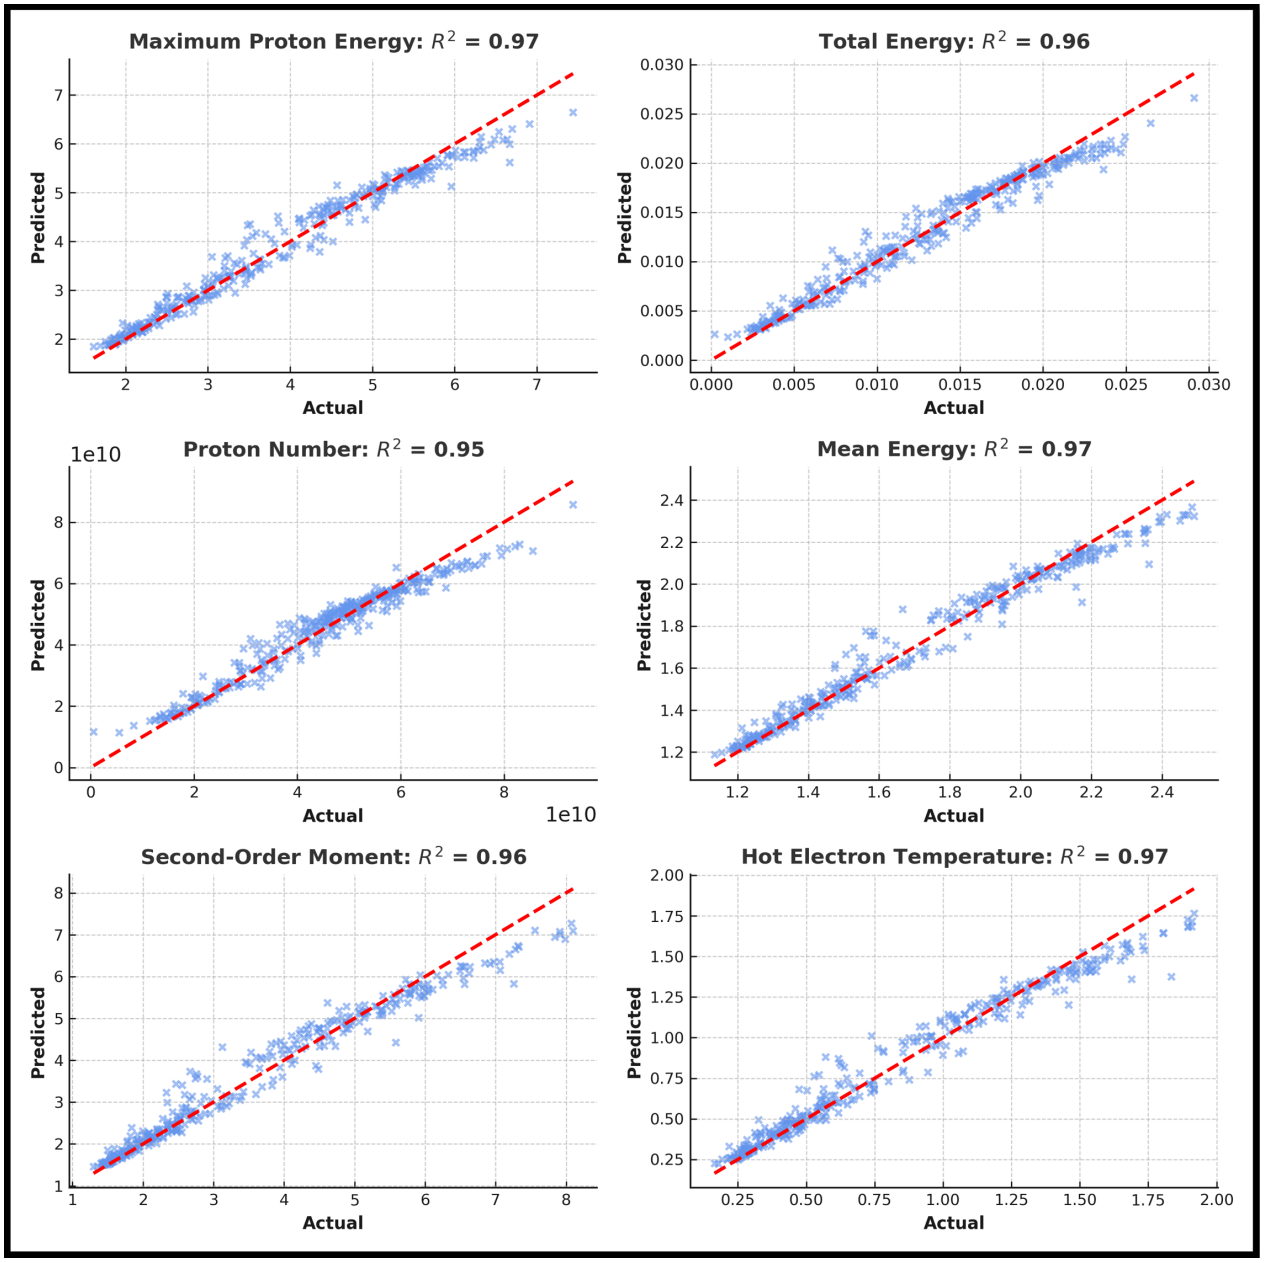

FIGURE 5.8 : RF predictions of proton beam metrics.

Prediction vs actual values for six key output metrics using RF regression models trained on DM actuator voltages. The dashed red line represents the ideal 1 :1 agreement. All models show excellent predictive performance, with coefficients of determination R^2 ranging from 0.95 to 0.97 for cutoff energy, total energy, proton number, mean energy, second-order momentum, and hot electron temperature. The slight deviation observed on the high side of the predicted parameter space likely originates from the reduced data density and higher experimental uncertainty in this region.

Figure 5.9 presents the distribution of prediction errors for the RF regression model trained to estimate proton maximum energies. The histogram represents the deviation between predicted and actual values ($\epsilon = y_{pred} - y_{true}$) across the entire dataset. The error distribution is approximately Gaussian and centered at zero, confirming that the model is unbiased and does not systematically over- or under-predict the output. Moreover, the majority of errors lie within ± 0.4 MeV. This narrow spread further supports the model's high precision and reinforces its suitability as a surrogate model in this study.

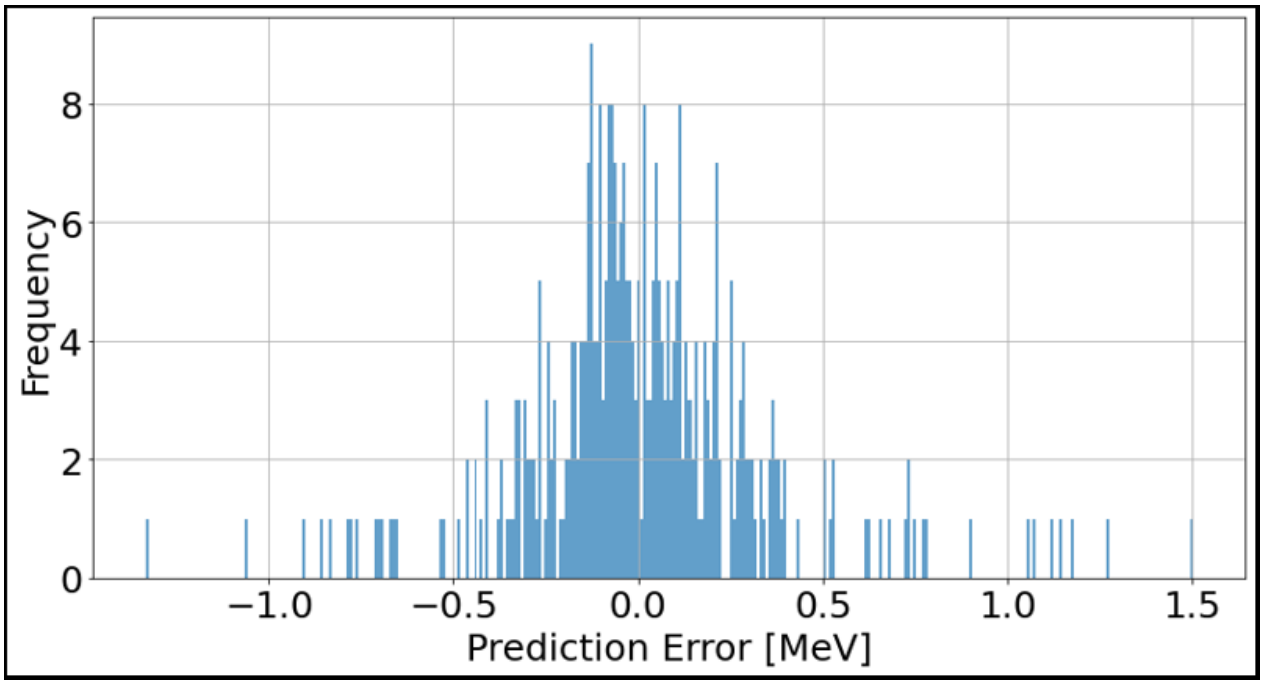

FIGURE 5.9 : Distribution of RF prediction errors for $\mathcal{E}_{p,max}$.

Distribution of prediction errors from the RF model trained to estimate proton maximum energy from DM actuator voltages. The histogram shows the difference between predicted and actual values across the full dataset. The distribution is centered near zero and approximately symmetric, indicating an unbiased model with low systematic error. Most errors fall within ± 0.4 MeV, reflecting the model's high accuracy and consistency.

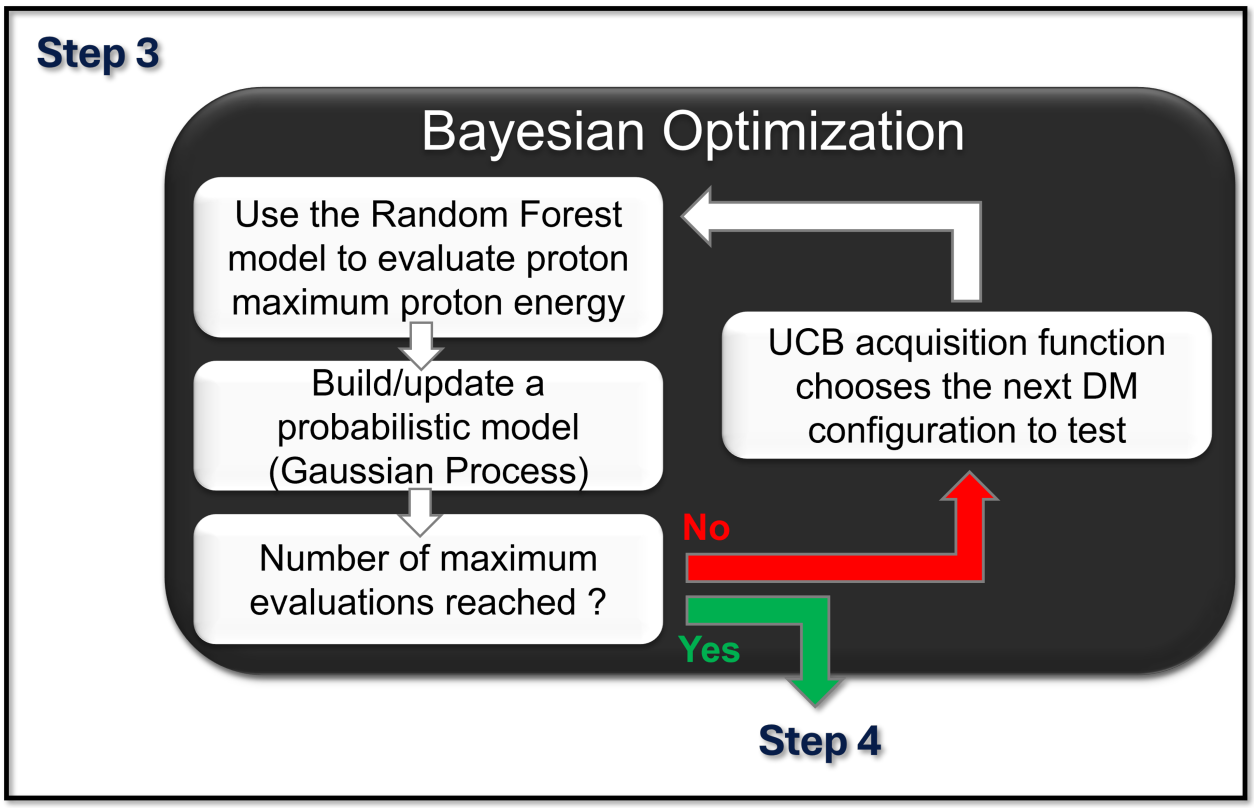

FIGURE 5.10 : Step 3 of the MSBO approach.

In the third step (Figure 5.10), BO was performed using Python's *gpminimize* function from *scikit-optimize* package (Contributors, 2025), with the trained RF model as the objective function. This GP-based method iteratively refined predictions by balancing exploration (sampling uncertain regions) and exploitation (focusing on promising configurations). The Upper Confidence Bound (UCB) acquisition function helped to efficiently navigate the high-dimensional, noisy search space, prioritizing high-energy or uncertain configurations. The Bayesian optimization was implemented using a Gaussian Process (GP) surrogate model with a Matérn 5/2 kernel, which provides a good compromise between smoothness and flexibility. This kernel assumes that the objective function is twice differentiable, allowing the model to capture moderately non-smooth behavior while avoiding overfitting to experimental noise. The kernel hyperparameters, including the characteristic length scale ℓ and the process variance σ_f^2, were optimized automatically at each iteration via marginal likelihood maximization. The acquisition function used was the Upper Confidence Bound (UCB) with $\kappa = 1$, balancing exploration and exploitation by sampling both high-uncertainty and high-predicted-yield regions of parameter space. A total of 120 iterations were performed, initialized with the zeroed DM actuator configuration. The optimization converged toward a stable optimum after approximately 70-80 evaluations, beyond which the acquisition function consistently proposed points within the same narrow region of the search space. To validate the robustness of the optimum, additional points were sampled in the vicinity of the predicted maximum, confirming that the model's response surface exhibited a well-defined local maximum rather than a noise-induced arti-

fact. This indicates that the identified optimum is not only numerically stable but also experimentally reproducible within the measurement uncertainty of the system.

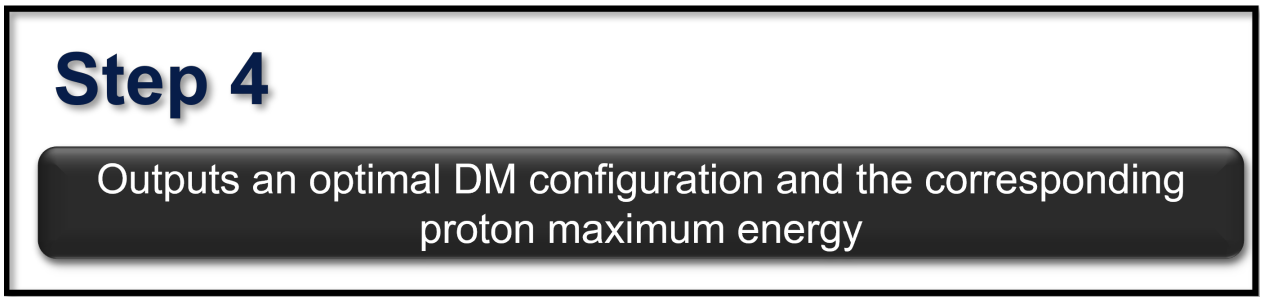

FIGURE 5.11 : Step 4 of the MSBO approach.

In the final step (Figure 5.11), BO identified an optimized configuration of the DM, and the corresponding maximum proton energy was predicted using the trained RF model. This configuration was subsequently tested experimentally to evaluate the performance of the optimization strategy.

Optimal values of a trained RF model can be found using random search, grid search, or optimization algorithms. BO was chosen for its efficiency in complex, high-dimensional spaces, guiding the search toward promising regions for maximizing proton energy. This iterative approach significantly reduces evaluations needed to converge to an optimal solution compared to an exhaustive random or grid search.

In low-repetition-rate experimental contexts, where each laser shot is time-consuming and resource-intensive, reducing the number of experimental evaluations is essential. BO offers an effective framework for this, even more so when paired with fast-to-evaluate surrogate models like RF. By using a trained RF model, BO can efficiently explore the search space without needing real experiments at each step, saving significant time and cost. The RF model handles noisy experimental data well and enables rapid evaluations, supporting robust offline optimization that smooths out noise and can ignore failed shots. It also reduces the risks of damaging configurations during live runs, like extreme voltages that could damage optics, by decoupling optimization from the experiment. This not only accelerates the search for optimal configurations but also facilitates safer and more flexible implementation. Our approach is well-aligned with established methods in the literature, such as Sequential Model-Based Algorithm Configuration (Hutter et al., 2011) and Efficient Global Optimization (Jones et al., 1998), which demonstrate the power of surrogate-assisted optimization in high-dimensional, costly-to-evaluate settings.

5.4 Results

A comparison of experimental proton spectra for three actuator configurations is shown in Figure 5.12. The maximum proton energy increased by 72%, reaching $\mathcal{E}_{p,max} = (6.64 \pm 0.30)$ MeV with the MSBO-determined optimal DM configuration. This represents a significant improvement over

the zeroed DM configuration, which yielded $\mathcal{E}_{p,max} = (3.85 \pm 0.14)$ MeV. MSBO also outperforms conventional wavefront correction methods, which optimizes the laser beam towards a flat wavefront, by ~24% ($\mathcal{E}_{p,max} = (5.37 \pm 0.24)$ MeV), demonstrating the superior performance of the MSBO-optimized configuration. Unlike MSBO, traditional wavefront correction lacks proton beam feedback when relying solely on the *SID4* WFS. $\mathcal{E}_{p,max}$ exhibits there a ~5% shot-to-shot fluctuation observed when the system is operated at a given actuator configuration, under nominally identical experimental conditions. This was assessed by analyzing 10 consecutive laser shots for fixed DM settings. Error bars in Figure 5.12 represent standard deviations over several repeated shots (10 per configuration). For the zeroed DM actuators, the WFS-corrected, and MSBO-optimized configurations, the relative standard deviations are :

$$\frac{0.14 \text{ MeV}}{3.85 \text{ MeV}} \approx 3.6\%, \quad \frac{0.24 \text{ MeV}}{5.37 \text{ MeV}} \approx 4.5\%, \quad \frac{0.30 \text{ MeV}}{6.64 \text{ MeV}} \approx 4.5\%$$

.

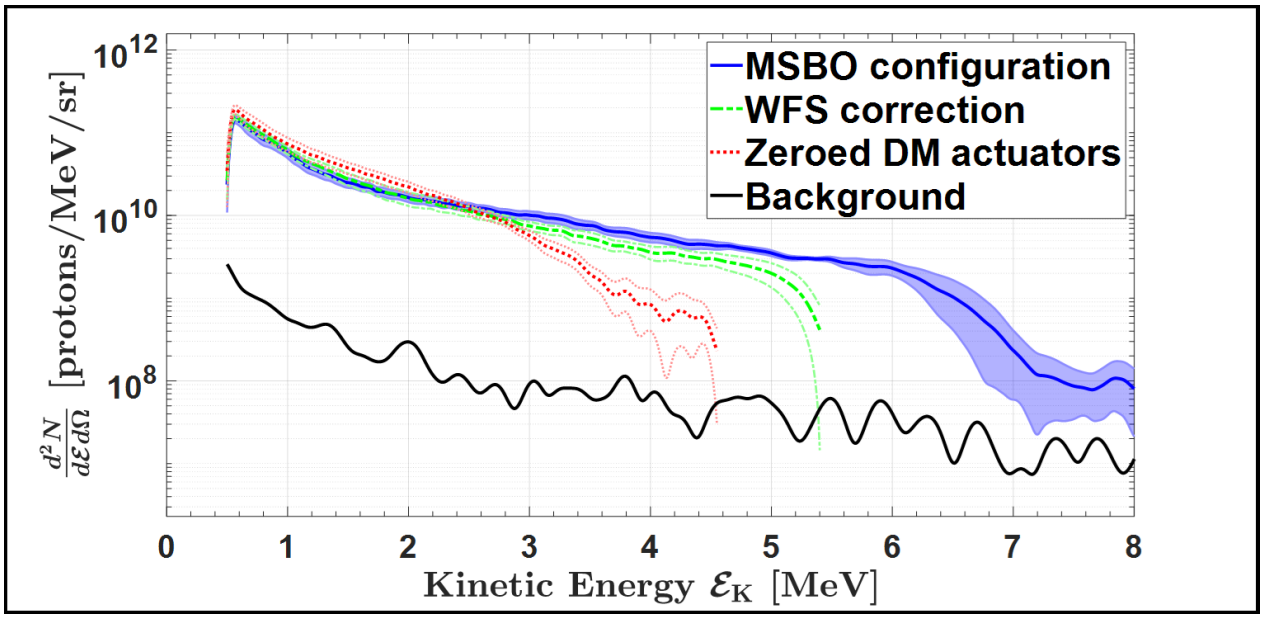

FIGURE 5.12 : Proton spectra for various DM configurations.

Average proton spectra over 10 shots under different optimization conditions; The window has been cut to 8 MeV for visualization purposes, but energies up to 12 MeV can be measured; Error bars are calculated using the standard deviation over 10 laser shots. This figure has been adapted from Catrix et al. (2025), with the publisher's permission.

The results reveal that the highest maximum proton energy consistently corresponds to a deformed wavefront, as shown in Figure 5.13. A-C and 5.14 .A-C, which present beam profiles without corrections, with WFS-based optimization, and with MSBO, respectively. Notably, maximum proton energy is achieved without maximizing laser intensity on the target. Figure 5.13.D highlights focal spot differences under identical conditions (same camera settings and gain). These findings

suggest that BO optimizes proton acceleration by shaping the wavefront for a more favorable laser-plasma interaction, even with lower peak intensity and broader width than WFS-based optimization.

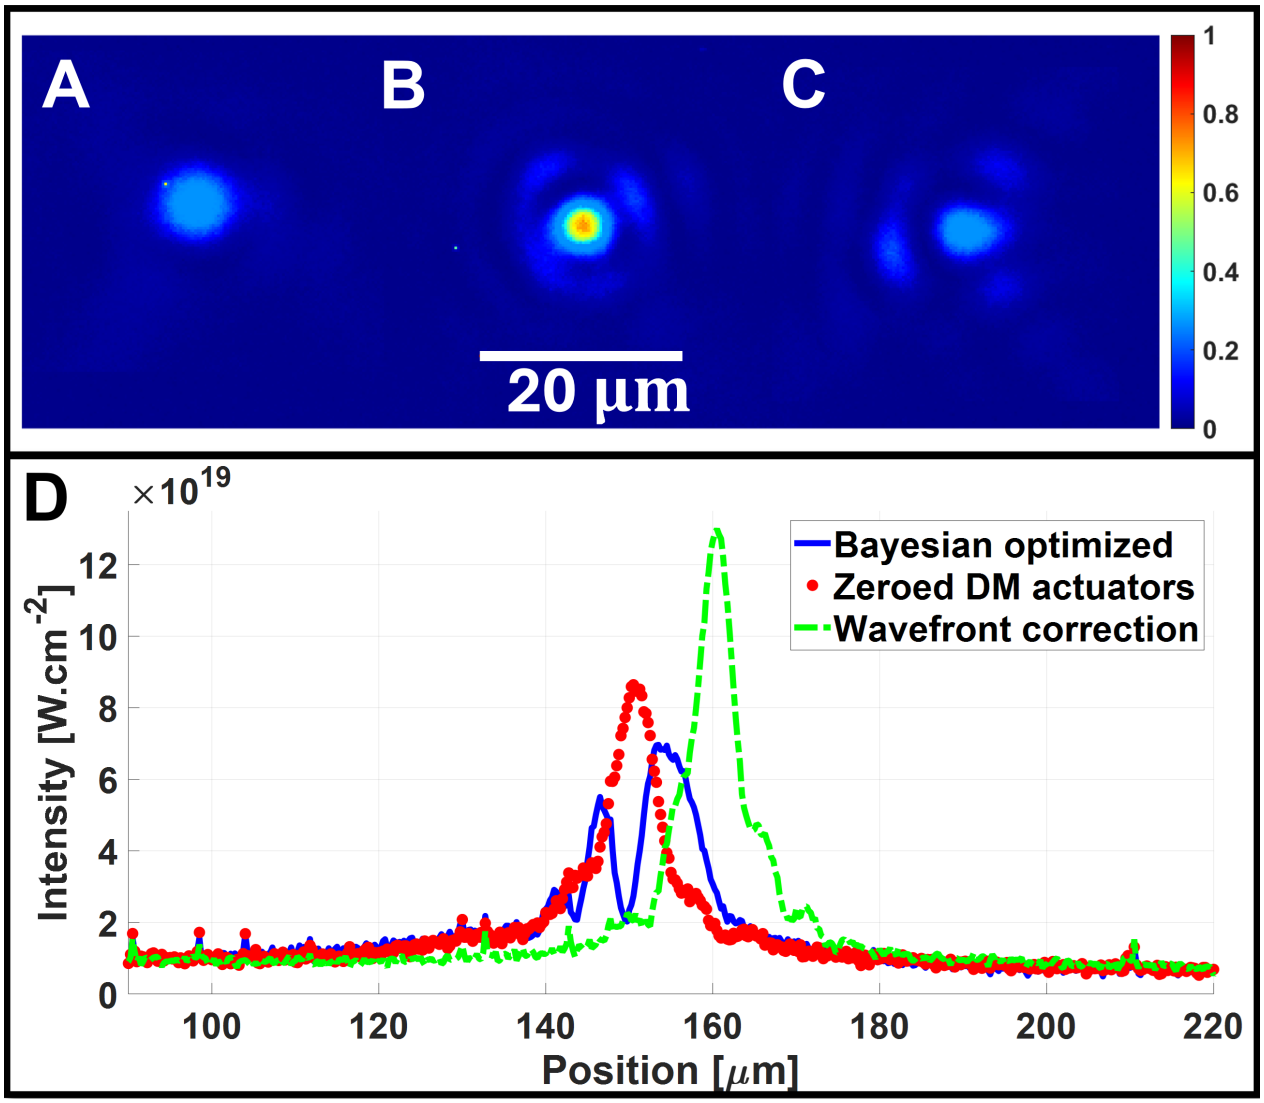

FIGURE 5.13 : Focal spots for various DM configurations.

Focal spots with (A) no wavefront correction, (B) optimized towards a flat wavefront using a WFS, and (C) optimized with MSBO. (D) Comparison of focal spot transverse (horizontal) profiles for 3 wavefront optimization strategies. The intensity is plotted as a function of position for three different cases : MSBO (solid blue), no correction (dotted red), and WFS correction (dashed green). These figures have been adapted from Catrix et al. (2025), with the publisher's permission.

5.5 Comparison with established TNSA scaling laws

To place the enhancement in proton acceleration achieved with our MSBO configuration into perspective, we compare our experimental results with established scaling laws from the literature.

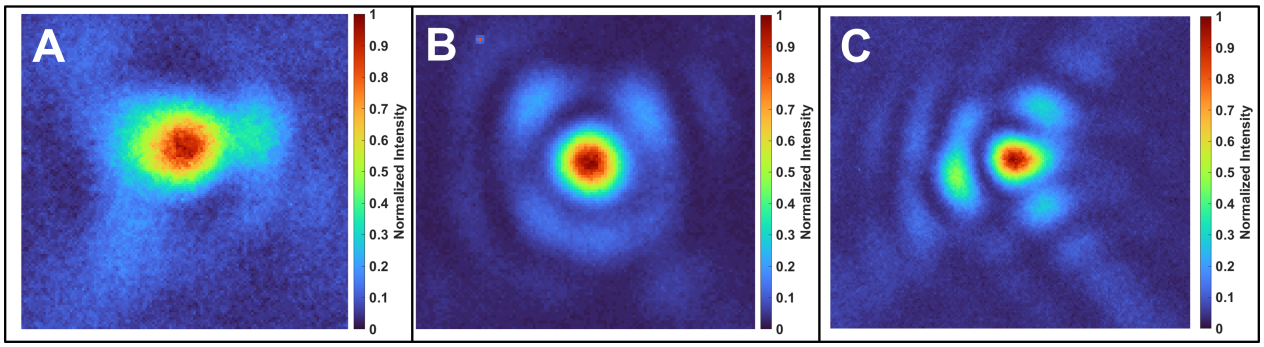

FIGURE 5.14 : Focal spots with (A) no wavefront correction, (B) optimized towards a flat wavefront using a WFS, and (C) optimized with MSBO.

Although based on different physical assumptions and valid in distinct regimes, these models provide a useful framework to interpret how laser–plasma interaction parameters govern $\mathcal{E}_{p,max}$ in the TNSA mechanism.

5.5.1 Fuchs et al. (2006) scaling model

Fuchs et al. (2006) model predicts the maximum proton energy based on the hot electron temperature (T_{hot}) and the normalized acceleration time (t_p), and is given by equation 1.30 using the estimation :

$$T_{hot} = m_e c^2 \left[\left(1 + \frac{I\lambda^2}{1.37 \times 10^{18}\,\mathrm{W \cdot cm^{-2} \cdot \mu m^2}} \right)^{1/2} - 1 \right], \tag{5.2}$$

where $t_p = \omega_{pi} \frac{t_{acc}}{\sqrt{2e}}$, with $\omega_{pi} = \sqrt{\frac{Z_i e^2 n_e^0}{m_i \epsilon_0}}$ being the ion plasma frequency and $t_{\mathrm{acc}} = 1.3\,\tau_{laser}$ the acceleration time. Here, e is the elementary charge, ϵ_0 is the vacuum permittivity, and T_{hot}, n_e^0 are the temperature and density of the hot electrons that drive the rear-surface expansion.

Although this scaling law is derived from a one-dimensional isothermal expansion model — which is a simplification of the actual three-dimensional plasma dynamics — it has been shown to accurately describe experimental results for TNSA-driven proton acceleration, particularly for laser pulse durations ranging from 40 fs to 10 ps, and intensities up to 6×10^{19} W.cm^{-2}, which is slightly different from our highest intense regime (22 fs, 1.3×10^{20} W.cm^{-2}). While the model is a simplified approximation to the actual 3D behavior, it remains a robust and widely used tool for interpreting proton energy scaling.

In our case, we find that despite a lower peak intensity in the MSBO condition, the measured hot electron temperature is higher ($T_{hot}^{MSBO} \approx 1.58$ MeV) than in the wavefront-corrected (WFS) case (≈ 1.16 MeV). This observation contrasts with the model's intensity-based prediction of T_{hot}. This result highlights that in TNSA the hot electron temperature is not determined solely by peak laser intensity, as assumed in the ponderomotive scaling of Equation 4.2. In our MSBO configuration, improved spatio-temporal wavefront shaping enhances laser absorption and energy transfer into

hot electrons, leading to a higher T_{hot} despite the lower on-target intensity. The stronger sheath field resulting from this hotter electron population directly contributes to the observed increase in $\mathcal{E}_{p,max}$. This underscores the limitation of one-dimensional scaling laws, which neglect multidimensional effects such as spatial focusing, plasma expansion geometry, and modified absorption channels.

5.5.2 Robson et al. (2007) empirical scaling

Robson et al. (2007) reported an empirical scaling law for the maximum proton energy based on an extensive dataset of PW-class laser–plasma experiments. Their study extended the investigated range of laser parameters up to 400 J of laser energy, intensities of 6×10^{20} W.cm^{-2}, and pulse durations up to 10 ps. They found that the increase of the maximum proton energy with laser intensity was significantly slower than previously predicted by isothermal expansion models such as Fuchs et al. (2006).

In particular, for a constant pulse duration of 1 ps and target thicknesses of 10–25 μm, the maximum proton energy $\mathcal{E}_{p,max}$ was shown to follow a simple power-law scaling expressed in equation 1.32. This dependence indicates that the maximum proton energy is proportional to the hot-electron temperature, consistent with ponderomotive scaling ($T_{hot} \propto (I\lambda^2)^{1/2}$). For intensities ranging from 4×10^{19} to 6×10^{20} W.cm^{-2}, $\mathcal{E}_{p,\max}$ increased only from about 10 MeV to 55 MeV.

They also investigated the dependence on pulse duration τ_L at fixed intensity ($\sim 8 \times 10^{19}$ W.cm^{-2}), finding that $\mathcal{E}_{p,max}$ increased only weakly, from 19 MeV at 1 ps to 24 MeV at 8 ps. This weak sensitivity contrasts with the predictions of Fuchs et al. (2006), who reported much stronger increases at lower intensities and shorter pulse durations.

Although the Robson scaling does not provide a closed analytical expression for $\mathcal{E}_{p,max}$ in terms of all laser parameters, it serves as an empirical benchmark showing that at high intensities ($> 10^{20}$ W.cm^{-2}), multidimensional effects and time-dependent electron dynamics significantly limit proton energies compared to 1D isothermal predictions. To capture this, Robson *et al.* extended the Mora isothermal expansion model (Mora, 2003) by introducing a two-phase description of the hot-electron temperature : a linear rise during the laser pulse followed by an adiabatic decay as in Mora (2005). This refinement accounts for the fact that when the electron temperature peaks, the rear surface already exhibits a finite density gradient, which weakens the sheath field and reduces acceleration efficiency. As a result, the predicted maximum energies are lower than the isothermal model, yet still about a factor of three higher than experimental measurements, underscoring the limitations of one-dimensional approaches in reproducing multidimensional laser–plasma coupling.

In our case, with ultrashort pulses of 22 fs on thin 4.5 µm aluminum foils, the MSBO configuration increases the hot-electron temperature and maximum proton energy beyond the values predicted by intensity-based scaling. This improvement stems from surpassing the performance of the WFS DM configuration, which provides a higher effective laser intensity through spatial mode control.

These results highlight that the dynamics are not governed solely by temporal heating and cooling, as assumed in the two-phase model, but also depend critically on spatial optimization. It is therefore important to note that the empirical Robson scaling derived for picosecond pulses (1–10 ps) and thicker targets (10–25 μm) cannot be directly extrapolated to our regime. This reinforces that intensity-based scaling is not suitable under our conditions and underscores the necessity of adaptive optimization strategies such as MSBO to capture the multidimensional physics at play.

5.5.3 Flippo et al. (2008) scaling at 200 TW

Flippo et al. (2008) provided another empirical scaling law from experiments on the *Trident* 200 TW laser, reporting in equation 1.33, with observations spanning 10^{18} to 10^{20} W.cm^{-2}. Their results are consistent with the relativistic ponderomotive scaling and support the notion that proton energy is not only a function of hot electron temperature and density, but also of sheath field geometry and temporal coupling. In their commissioning of the 200 TW *Trident* laser, Flippo *et al.* reported that optimized laser conditions, particularly high contrast and improved focusing, allowed the system to outperform empirical scaling curves at sub-10^{19} $\mathrm{W.cm^{-2}}$ intensities, although they noted that *Trident* comes out ahead of the curves for intensities below 1×10^{19} $\mathrm{W.cm^{-2}}$, but falls along the scaling laws above that value.

Our observations are consistent with the findings of Flippo et al. (2008), who stressed that determined scaling laws are apparently laser dependent due to some of the other individual laser characteristics that affect the beam quality and beam energy, such as the laser focal spot, prepulse contrast, and ASE contrast. However, at 7×10^{19} $\mathrm{W.cm^{-2}}$ we fall in a regime where Flippo *et al.* observed convergence back to intensity-based scaling. In this context, *Trident*-type scaling is not completely fitted to our case and validate the use of adaptive strategies such as MSBO to exploit degrees of freedom.

5.5.4 Zeil et al. (2010) linear scaling with ultrashort pulses

Zeil et al. (2010) investigated TNSA with 30 fs, 100–150 TW Ti : Sapphire pulses and few-µm foils, reporting maximum proton energies up to 17 MeV on *Draco* laser and, crucially, a departure from the traditional $\sqrt{P_L}$ law toward a linear scaling with laser power in the ultrashort regime. They interpret the data with the Schreiber surface-charge model (Schreiber et al., 2006), which yields

Equation 5.2 : Asymptotic proton maximum energy in the ultrashort-pulse regime (Zeil et al., 2010)

$$E_\infty = 2m_e c^2 \sqrt{\frac{\eta P_L}{P_e}}, \qquad P_e \equiv \frac{m_e c^3}{r_e} \simeq 8.7 \text{ GW} \tag{5.3}$$

and

Equation 5.3 : Proton maximum energy in the ultrashort-pulse regime (Zeil et al., 2010)

$$\mathcal{E}_{p,\max} = E_\infty \, \tanh^2\!\left(\frac{\tau_L}{2\tau_0}\right), \qquad \tau_0 = \frac{R}{v_\infty}, \qquad R = r_f + d\,\tan\vartheta \tag{5.4}$$

Here η is the laser to electron conversion, P_L the laser power, τ_L the pulse duration, r_f the focal radius on target, d the foil thickness, and ϑ the hot-electron half-divergence inside the foil. Two asymptotic scalings follow : for $\tau_L \ll 2\tau_0$, $\mathcal{E}_{p,\max} \approx E_\infty(\frac{\tau_L}{2\tau_0})^2 \propto \eta P_L$ (linear in power), while for $\tau_L \gg 2\tau_0$, $\mathcal{E}_{p,\max} \to E_\infty \propto \sqrt{\eta P_L}$ (square-root regime). With *Draco* parameters (∼150 TW, ∼30 fs, $r_f \sim 1.7$ μm, $d = 2$–5 μm, $\vartheta \sim 10°$), the experiment falls in the short-acceleration-time regime and indeed exhibits near-linear power scaling up to ∼17 MeV. Rear-surface geometry also matters : preplasma-driven shocks can deform thin foils and steer the most energetic protons off the target normal, complicating spectral inference and indicating sensitivity to R through d, r_f and ϑ.

Interestingly, in our MSBO configuration, the wavefront shaping leads to an increase in focal spot size, which by itself would reduce the maximum proton energy within the Zeil framework (through a larger effective source size R). However, this effect seems to be more than compensated by an increase of the coupling efficiency η, together with enhanced sheath symmetry.

5.6 Laguerre–Gaussian mode analysis of the MSBO laser beam

To quantify the nature of the MSBO laser beam profile, we performed a Laguerre–Gaussian (LG) mode decomposition of the far-field intensity distribution after optimization. The decomposition was carried out in cylindrical coordinates using a complete orthonormal LG basis (LG_p^l) up to radial order $p = 5$ and azimuthal order $\ell = \pm 5$. The analysis (Figures 5.15 and 5.16) reveals that the optimized beam exhibits a complex structure composed of both low- and higher-order modes. The fundamental Gaussian mode LG_0^0 contributes about 29% of the energy, but substantial fractions are found in higher-order radial modes such as LG_2^0, LG_3^0, and LG_4^0. This demonstrates that the MSBO primarily reshapes the wavefront by introducing structured radial content, consistent with the nontrivial phase and intensity modulations observed in Figure 5.13.D.

Figure 5.15 provides the detailed modal decomposition across the full LG basis, while Figure 5.16 summarizes the energy grouped by radial (p) and azimuthal ($|\ell|$) indices. figure 5.16.A highlights the significant redistribution of energy across multiple radial orders, confirming the role of the DM in generating complex radial phase variations. Figure 5.16, in contrast, shows that more than 90% of the beam energy remains concentrated in azimuthally symmetric modes ($|\ell| = 0$), with only ∼10% carried by higher-order orbital angular momentum (OAM) modes. This indicates that the beam, despite being far from a pure Gaussian or low-order LG mode, retains strong cylindrical symmetry.

Such a configuration is particularly relevant in the TNSA context, as beams carrying OAM have been shown to enhance proton energies by modifying laser–plasma coupling in the near-critical layer (Willim et al., 2023). Indeed, they show that reduced relativistic self-focusing of OAM beams prevents premature relativistic transparency of the overdense foil, preserving a cold rear surface and maintaining cylindrically symmetric sheath fields. These symmetric fields confine hot electrons axially and minimize transverse work on the most energetic protons, thereby reducing lateral energy losses and allowing a larger fraction of the absorbed laser energy to be converted into forward proton momentum. This mechanism may allow, in our case, the MSBO laser beam to deliver proton beams with higher proton maximum energies than Gaussian drivers at the same laser energy, while simultaneously achieving much lower divergence.

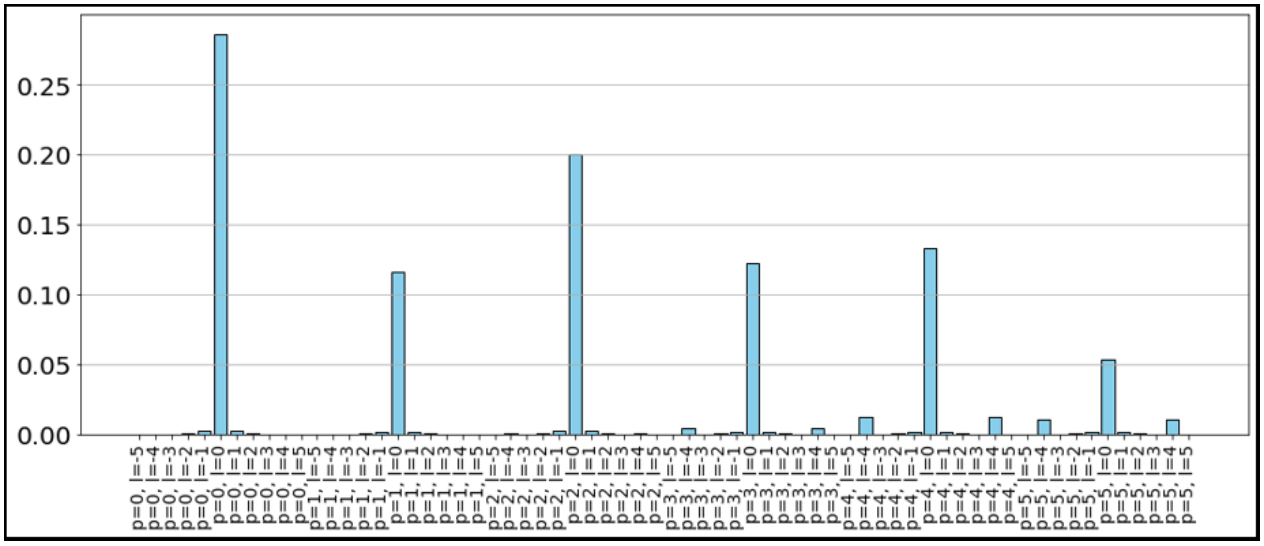

FIGURE 5.15 : Radial and azimuthal LG energy distribution of the MSBO laser beam.

Laguerre–Gaussian mode decomposition of the laser beam obtained under the MSBO configuration. The bar heights represent the normalized contribution of each LG mode, indexed by radial number p and azimuthal index ℓ, to the reconstructed beam profile.

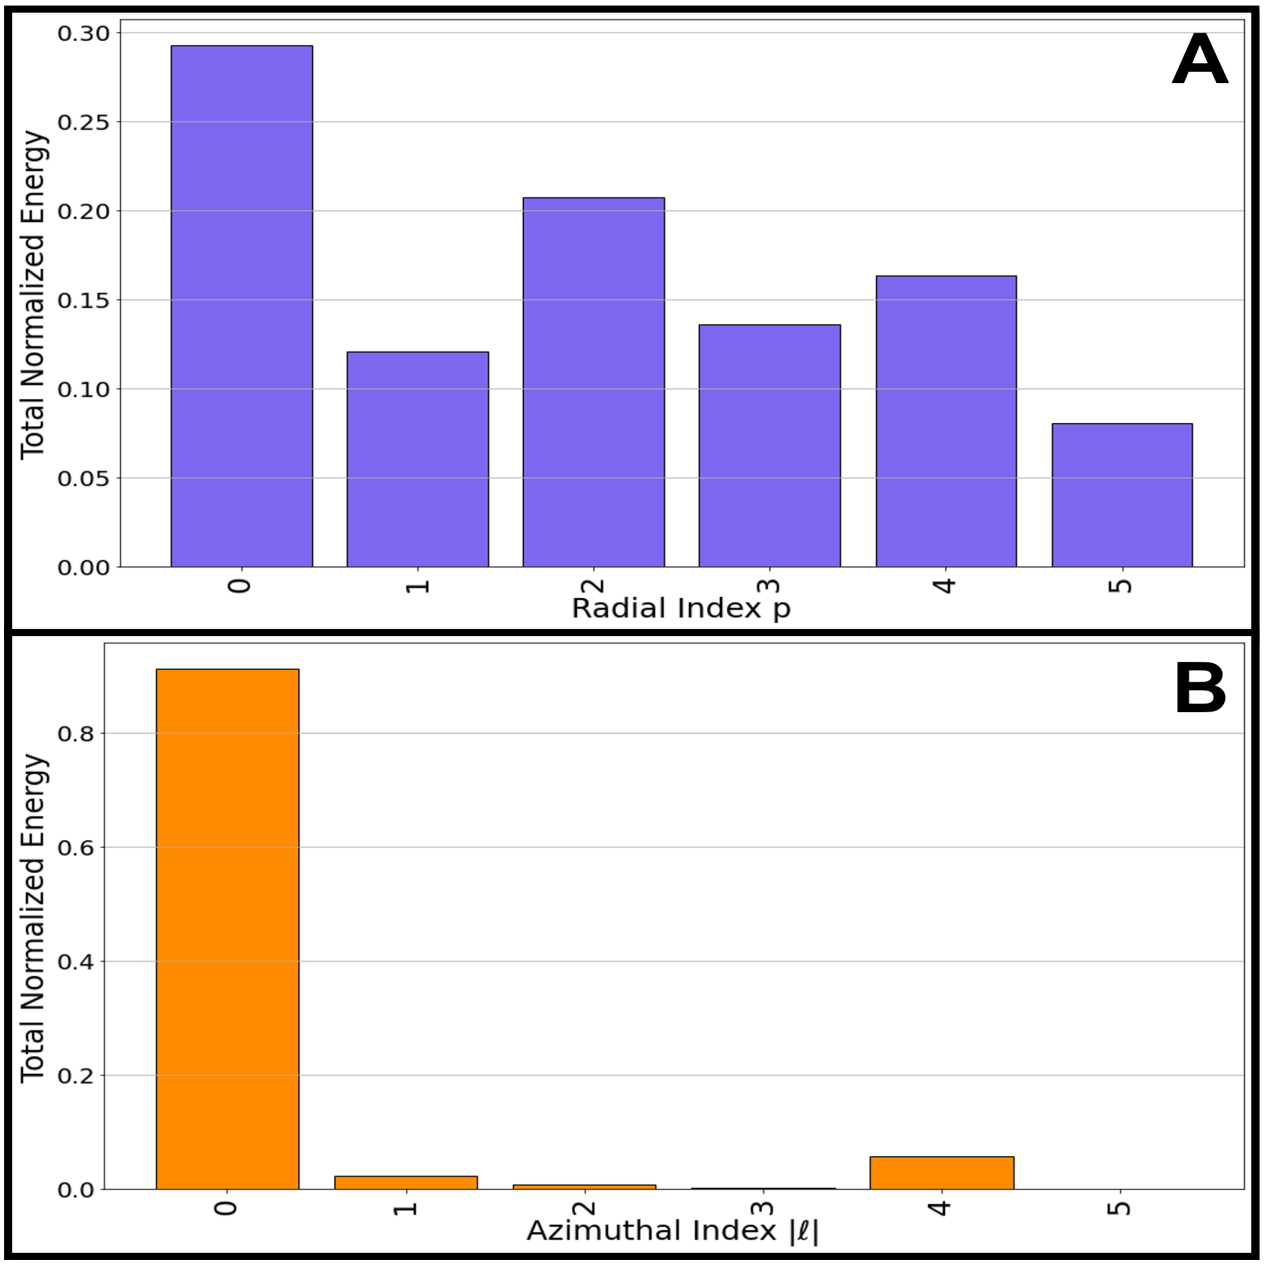

FIGURE 5.16 : LG mode decomposition of the MSBO-optimized laser beam.

(A) Total normalized energy per radial index p and (B) per absolute azimuthal index $|\ell|$ for the Laguerre–Gaussian decomposition of the MSBO laser beam.

This complex but structured profile suggests that, although the beam deviates significantly from a Gaussian, its energy may still be concentrated within a limited subset of modes. To better quantify this property, we analyze its decomposition into LG modes, focusing on the distribution of energy among the modes. The LG mode decomposition reveals that the MSBO-optimized beam, though non-Gaussian, can be reconstructed using only a few dominant modes, highlighting its compact structure and enabling simplified modeling of the laser–plasma interaction if required. Figure 5.17 presents the cumulative normalized energy distribution across the first 66 modes in the LG decomposition of the MSBO-optimized beam. The curve illustrates how many modes are required to capture a given fraction of the total beam energy. The sharp initial rise and early saturation reveal that 95% of the beam's energy is contained within the first 10 modes, with higher-order modes contributing negligibly. This demonstrates that, despite its non-Gaussian profile, the optimized beam can be effectively represented using a low-dimensional LG mode set. Such modal compactness is valuable for both physical understanding and further computational modeling for precisely understanding the underlying processes that are happening through the interaction of our MSBO beam with solid targets.

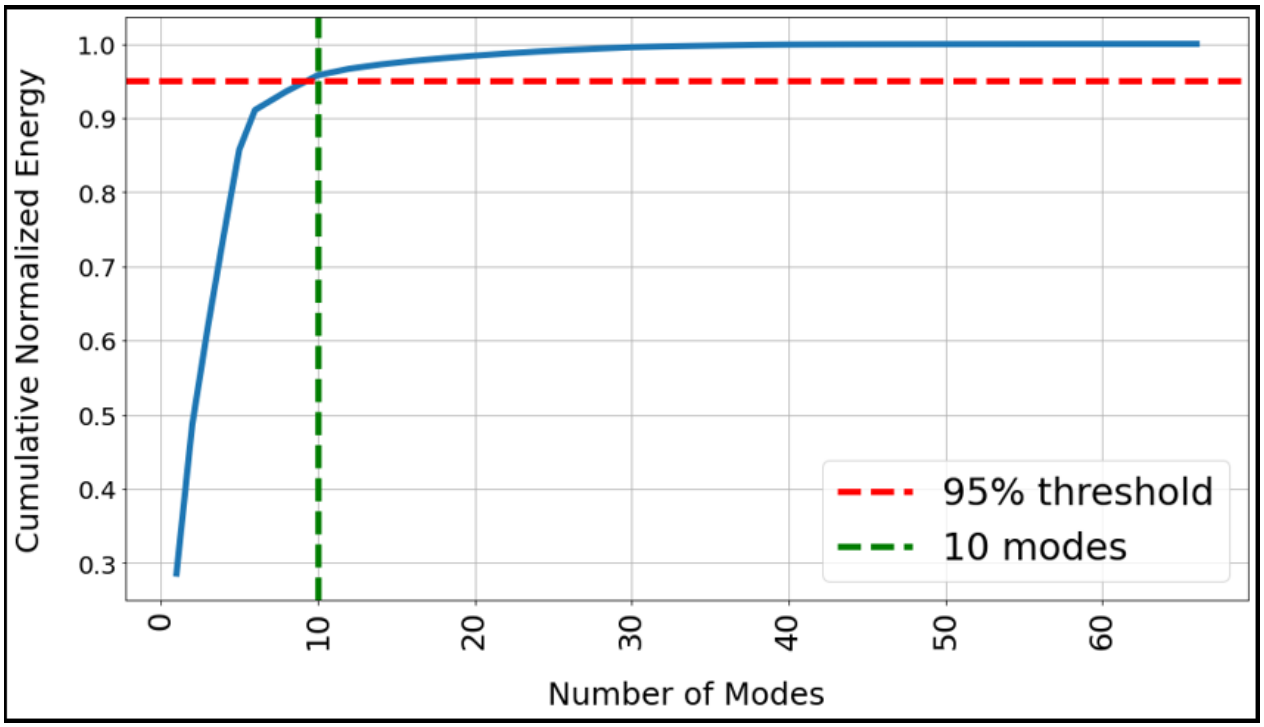

FIGURE 5.17 : Cumulative LG energy vs number of modes for MSBO beam reconstruction

Cumulative normalized energy as a function of the number of LG modes included in the reconstruction of the MSBO beam. The red dashed line marks the 95% energy threshold, which is reached with only the first 10 modes (vertical green dashed line). This demonstrates that the beam profile can be effectively represented using a compact subset of the full LG basis.

5.7 Convergence and predictive accuracy of data-driven MSBO

To evaluate the efficiency and reliability of the optimization process, we conducted 89 consecutive optimization trials, each one using a different RF predictive model trained on a randomly selected and size-increased subset of 2 to 358 laser shots from the initial dataset (2, 6, 10, 14 ... up to 358). Each optimization trial was repeated 10 times for statistical robustness, tracking the global maximum proton energy predicted by the GP model at each iteration. All runs started with DM actuator voltages set to 0 V. Figure 5.18.A shows the mean (blue curve) and standard deviation (shaded sky blue curve) of the predicted optima across ten runs, illustrating the evolution of $\mathcal{E}_{p,max}$ as a function of the number of laser shots. Initially, $\mathcal{E}_{p,max}$ is $\sim$3 MeV for $\sim$5 shots, rising sharply to $\sim$90% of its final value within 30 shots and reaching 6 MeV around 50 shots. The growth then slows down, stabilizing around 6.2 MeV after about 100 shots, with minor fluctuations up to 358 shots. This suggests that 100-150 shots suffice to determine an optimal DM configuration and reliably estimate $\mathcal{E}_{p,max} = (6.44 \pm 0.33)$ MeV, closely matching the experimental value ($\mathcal{E}_{p,max} = (6.64 \pm 0.30)$ MeV). The fitting function (orange curve) is defined as :

Equation 5.4 : Empirical convergence law under MSBO for the maximum proton energy

$$\mathcal{E}_{p,max}(N_{\text{shots}}) = 3.337\left(1 - e^{-N_{\text{shots}}/12.23}\right) + 2.948 \tag{5.5}$$

where N_{shots} is the number of laser shots. The shaded region represents uncertainty, initially higher due to limited data but narrowing as more shots improve the model's confidence. This analysis highlights the trade-off between data acquisition and predictive accuracy, crucial for optimizing experiments. The entire process, including dataset collection, parameter setting, and computation (5 minutes), took $\sim$1 hour. In comparison, manual and wavefront correction methods require at least a similar timeframe but yield suboptimal results. Optimizing only Actuator0, which correlates most with $\mathcal{E}_{p,max}$, resulted in $\mathcal{E}_{p,max} = (5.25 \pm 0.22)$ MeV, demonstrating the advantage of high-dimensional wavefront control in laser-plasma acceleration.

In Figure 5.18.B, we represent the evolution of the RMSE of RF predictive models as a function of the number of laser shots used to train and test the model. Here, RMSE is computed on experimental test datasets (15% of the initial data) to evaluate model error and optimization convergence. Initially high ($\sim$1.3 MeV for <10 shots) with large fluctuations due to limited training data, RMSE decreases rapidly, falling below 0.5 MeV around 50 shots. Beyond $\sim$100 shots, it stabilizes at 0.25 to 0.35 MeV, with minor fluctuations up to 358 shots, indicating that additional data beyond this threshold yield minimal accuracy gains. The residual noise ($\sim \pm 0.3$ MeV) at higher shot numbers likely stems from inherent experimental measurement noise, in particular shot-to-shot fluctuations.

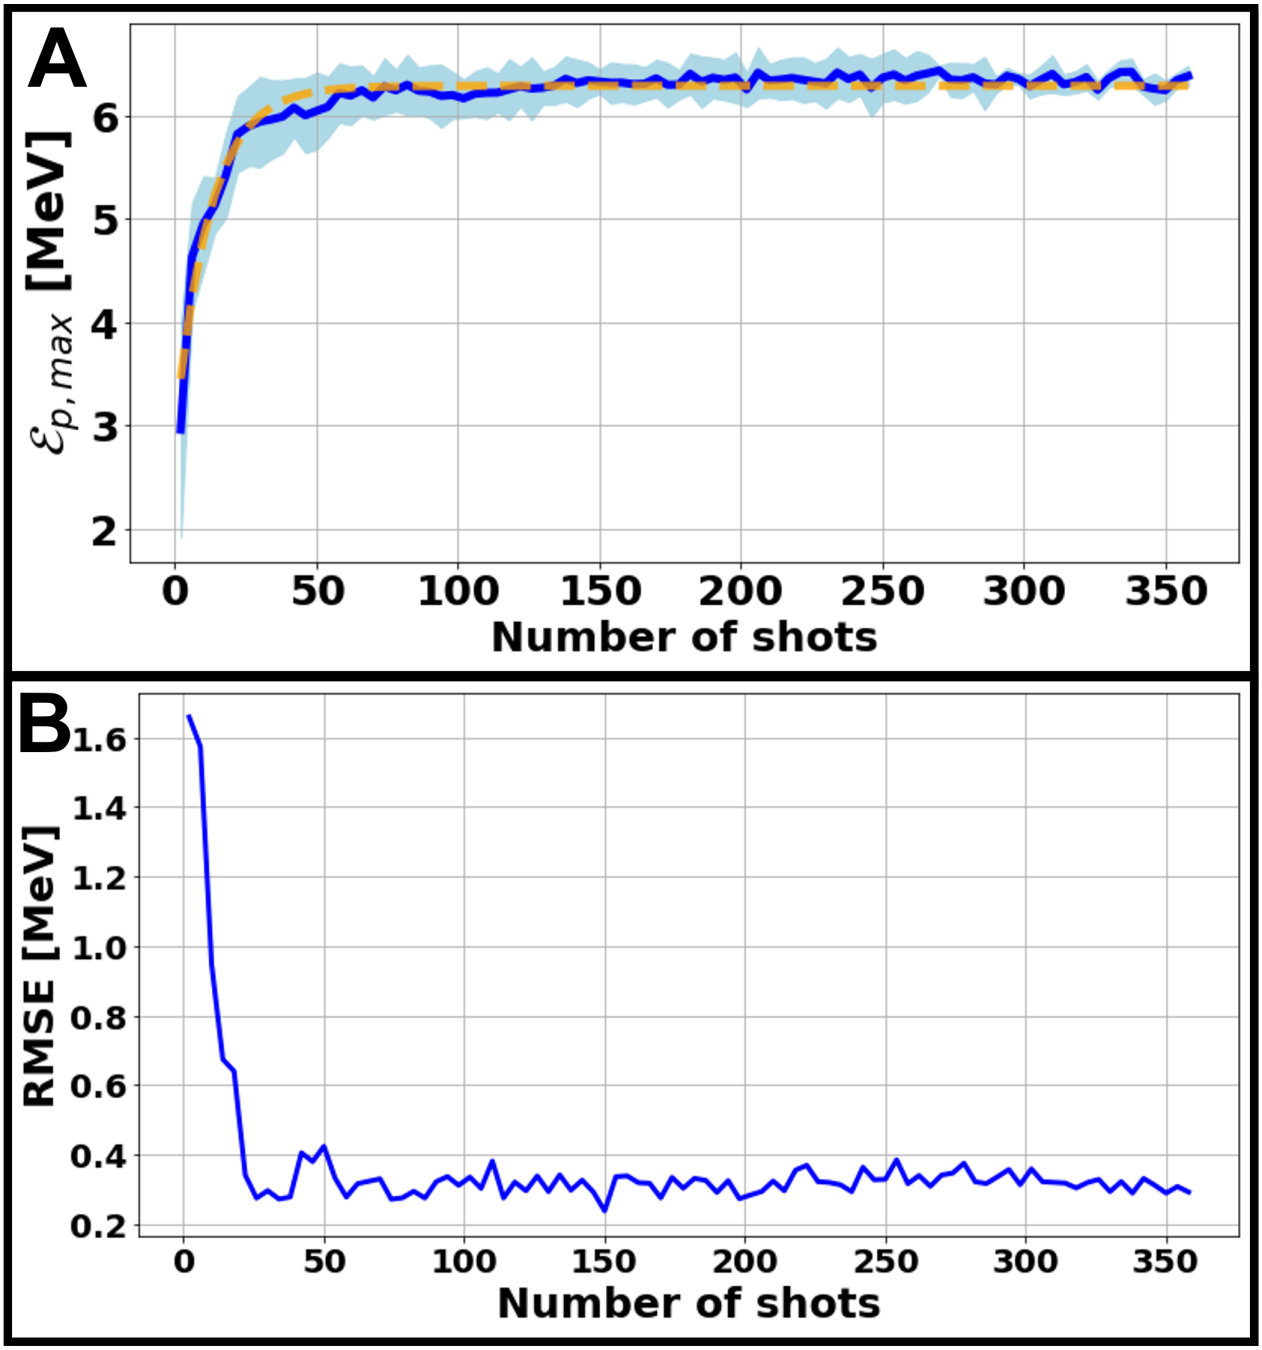

FIGURE 5.18 : Convergence of $\mathcal{E}_{p,max}$ prediction and RF model RMSE.

(A) Average convergence plot for maximum estimated proton energy selecting successively and randomly 1 to 360 data samples to build the predictive model, over 10 runs of MSBO. (B) Evolution of the Root Mean Square Error of the RF predictive model as a function of shots used to train and test the model. These figures have been adapted from Catrix et al. (2025), with the publisher's permission.

5.8 Zernike decomposition and wavefront analysis

To assess the optical quality and modal structure of the laser beam, we performed a Zernike decomposition of the wavefronts measured using the *SID4* WFS. This decomposition quantifies the contribution of optical aberrations (e.g., defocus, astigmatism, coma) to the overall wavefront error, providing direct insight into how each DM configuration affects the laser beam quality and, consequently, the laser–plasma coupling efficiency. Three distinct configurations of the DM were analyzed :

— (A) MSBO-optimized

— (B) WFS-based correction

— (C) Zeroed DM actuators

All configurations were analyzed using the first 21 Zernike modes.

5.9 Zernike polynomials and their mathematical representation

Zernike polynomials are orthogonal functions defined on the unit disk, used extensively in optics to describe wavefront aberrations (Niu et al., 2022). Because Zernike polynomials are orthogonal on the unit disk and correspond naturally to classical aberration modes, they provide a mathematically rigorous and physically intuitive basis for analyzing, correcting, and optimizing optical wavefronts. Their use in adaptive optics and wavefront sensing enables precise correction of aberrations, improving image quality and system performance. Each polynomial is described by a radial order $n \geq 0$ and azimuthal frequency m, where $|m| \leq n$ and $n - |m|$ is even.

The general form is :

Equation 5.5 : General definition of Zernike polynomials on the unit disk

$$Z_n^m(r,\theta) = \begin{cases} R_n^m(r)\,\cos(m\theta), & m > 0, \\ R_n^{|m|}(r)\,\sin(|m|\theta), & m < 0, \\ R_n^0(r), & m = 0, \end{cases} \tag{5.6}$$

with the radial component defined as :

Equation 5.6 : Radial component of the Zernike polynomial

$$R_n^m(r) = \sum_{k=0}^{\frac{n-m}{2}} (-1)^k \, \frac{(n-k)!}{k!\,(\frac{n+m}{2}-k)!\,(\frac{n-m}{2}-k)!}\, r^{\,n-2k}. \tag{5.7}$$

5.10 Zernike indexing conventions

5.10.1 Noll indexing scheme

To simplify the handling and referencing of these modes in optical metrology and computational tools, R.J. Noll introduced a one-dimensional indexing scheme (Noll, 1976) that maps each (n, m) pair to a single scalar index j.

The so-called *Noll index* assigns a sequential number to Zernike modes in a specific order. This scalar indexing is especially useful in wavefront sensing and reconstruction. Table 5.2 summarizes the first 21 Zernike polynomials according to Noll's indexing.

TABLE 5.2 : First 21 Zernike modes using Noll's convention.

Noll Index j	Zernike Mode	Radial Order n	Azimuthal Order m
1	Piston	0	0
2	Tilt X	1	1
3	Tilt Y	1	-1
4	Defocus	2	0
5	Astigmatism @ 0°	2	2
6	Astigmatism @ 45°	2	-2
7	Coma X	3	1
8	Coma Y	3	-1
9	Trefoil @ 0°	3	3
10	Trefoil @ 30°	3	-3
11	Spherical Aberration	4	0
12	Secondary Astigmatism @ 0°	4	2
13	Secondary Astigmatism @ 45°	4	-2
14	Tetrafoil @ 0°	4	4
15	Tetrafoil @ 25.5°	4	-4
16	Second Coma X	5	1
17	Second Coma Y	5	-1
18	Second Trefoil @ 0°	5	3
19	Second Trefoil @ 30°	5	-3
20	Pentafoil @ 0°	5	5
21	Pentafoil @ 18°	5	-5

Each Zernike mode corresponds to a distinct type of optical aberration that is commonly encountered in laser systems and adaptive optics :

— $j = 2, 3$: tip and tilt (beam pointing errors)

— $j = 4$: defocus (curvature, convergence/divergence)

— $j = 5, 6$: astigmatism (elliptical deformation)

— $j = 7, 8$: coma (off-axis aberration)

— $j = 9, 10$: trefoil (triangular distortion)

— $j = 11$: spherical aberration (peripheral focus shift)

The phase $\Phi(x, y)$ is reconstructed as :

Equation 5.7 : Wavefront reconstruction using Noll-indexed Zernike polynomials

$$\Phi(x, y) = \sum_{j=1}^{N} a_j \, Z_j(x, y), \tag{5.8}$$

where Z_j is the Noll-indexed Zernike polynomial and a_j is its corresponding coefficient.

5.10.2 *Phasics SID4* convention

The *Phasics SID4* WFS uses a decomposition that retains the native (n, m) structure and outputs two separate coefficients :

— $c^e_{n,m}$: coefficient for $\cos(m\theta)$ (even term, corresponds to $+m$)

— $c^o_{n,m}$: coefficient for $\sin(m\theta)$ (odd term, corresponds to $-m$)

Only modes with $m \geq 0$ are explicitly listed, with both signs of m encoded via trigonometric components. The phase is reconstructed as :

Equation 5.8 : Phase reconstruction using the *Phasics SID4* convention

$$\Phi(x, y) = \sum_{n,m \geq 0} \left[c^e_{n,m} \, R^m_n(r) \cos(m\theta) + c^o_{n,m} \, R^m_n(r) \sin(m\theta) \right]. \tag{5.9}$$

This decomposition avoids redundancy and simplifies the interpretation of individual modal contributions. It is especially useful when analyzing experimental wavefronts captured by the *SID4* WFS.

While both methods rely on the same underlying Zernike basis, their representations differ :

TABLE 5.3 : Comparison of Noll and *Phasics* Zernike decompositions.

Aspect	**Noll Convention**	**Phasics Convention**
Indexing	Scalar index j	Explicit (n, m) pairs (with $m \geq 0$)
Angular terms	Combined sign in Z_j	Split into components $\cos(m\theta)$ and $\sin(m\theta)$
Coefficient	One per mode : a_j	Two per mode : $c^e_{n,m}$ (even), $c^o_{n,m}$ (odd)
Normalization	Orthonormal (optional RMS)	RMS-normalized by default
Interpretability	Compact scalar index	Explicit physical meaning for each mode

5.11 Wavefront reconstruction

Wavefront reconstruction from Zernike decomposition provides a quantitative and interpretable representation of the optical phase aberrations present in the beam. Despite the differences in representation, both conventions (Noll and *Phasics*) produce mathematically equivalent reconstructions of the wavefront. The consistency between them was verified by reconstructing the phase maps using both sets of coefficients and comparing the resulting phase maps.

Figure 5.19 shows the reconstructed wavefronts for the three DM configurations in Noll's convention.

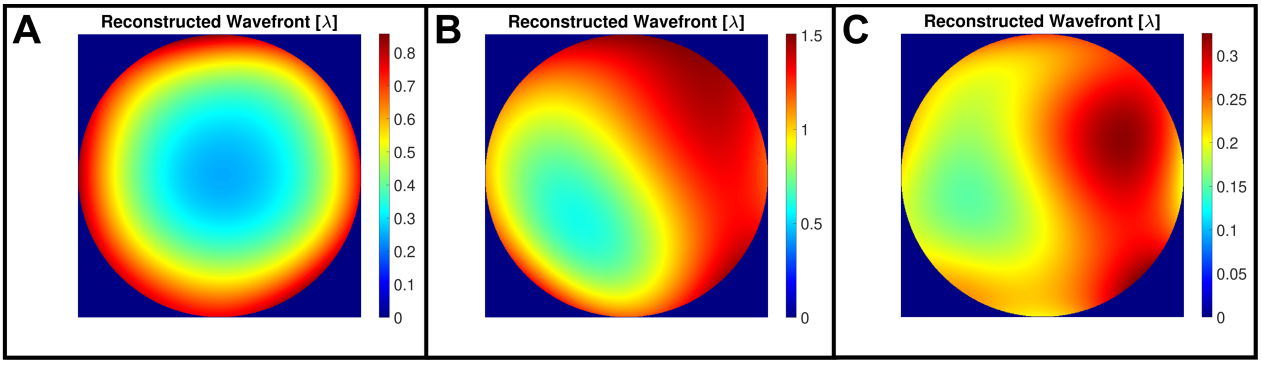

FIGURE 5.19 : Reconstructed wavefront maps for three DM configurations.

Reconstructed phase maps in units of λ for the three DM configurations in Noll's convention : (A) Zeroed DM actuators configuration, (B) MSBO DM configuration and (C) WFS-based optimization.

Figure 5.20 shows the reconstructed wavefronts for the three DM configurations in the *SID4 Phasics* convention. The phase map is flipped in this convention compared to Noll's convention.

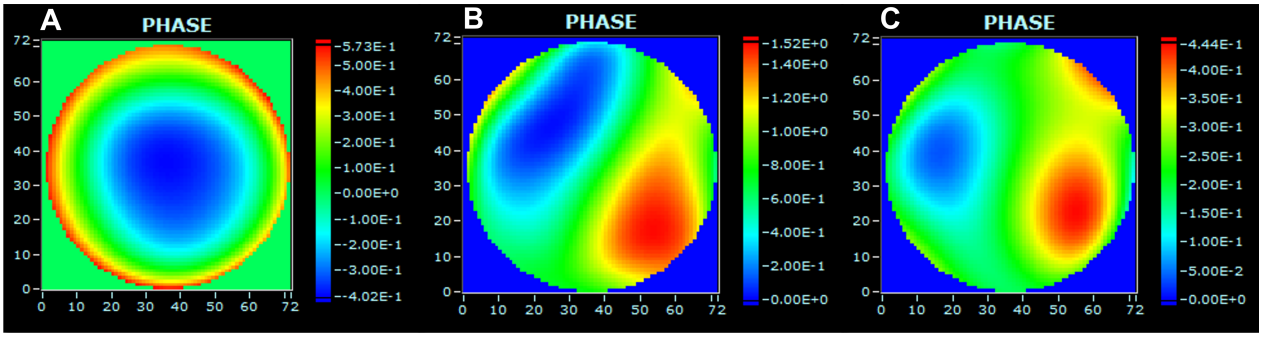

FIGURE 5.20 : *SID4* phase maps for three DM configurations.

Reconstructed phase maps in units of λ for the three DM configurations in the *SID4 Phasics* convention : (A) Zeroed DM actuators configuration, (B) MSBO DM configuration and (C) WFS-based optimization.

The *Phasics* analysis software shows that WFS correction yields the best beam quality in classical optical terms (Table 5.4), while the MSBO configuration deliberately accepts greater wavefront dis-

TABLE 5.4 : Wavefront quality metrics for three DM configurations.

Configuration	PtV Phase (λ)	RMS Phase (λ)	Strehl Ratio
Zeroed DM actuators	0.97	0.287	0.066
WFS	0.44	0.110	0.692
MSBO	1.52	0.432	0.179

tortion to achieve better outcomes. More precisely, in the zeroed DM actuators' configuration, the dominant aberration is defocus (0.265λ), accompanied by minor contributions from low-order astigmatism (e.g., -0.012λ and 0.016λ) and coma (0.004λ, -0.008λ). This suggests a relatively simple but uncorrected wavefront error typical of a relaxed or initial DM state. In the WFS-corrected configuration, all aberrations are significantly reduced, with defocus (0.009λ) and coma X (-0.056λ) being the largest residual components, indicating that the WFS-based correction effectively minimizes both low- and mid-order terms, resulting in the highest Strehl ratio. In contrast, the MSBO configuration presents a markedly different aberration profile, with significant contributions from coma Y (-0.189λ), tilt Y (-0.185λ), tilt X (0.272λ), and astigmatism (0.124λ). Despite the resulting lower Strehl ratio (0.179), this configuration yields to a better TNSA efficiency. These results suggest that the MSBO strategy exploits deliberate low- and mid-order aberrations, particularly coma and astigmatism, to enhance laser energy deposition and hot electron generation, ultimately improving ion acceleration efficiency in the TNSA regime.

5.11.1 Energy spectrum per configuration

In Noll's convention, the phase $\Phi(x, y)$ over the unit pupil is reconstructed using equation 5.7 with N=21. Modal energy is given by :

Equation 5.9 : Modal energy in Noll's convention

$$E_j = a_j^2 \tag{5.10}$$

We use here the Noll-indexed decomposition to compute modal energy spectra. Figure 5.21 shows the Zernike energy modal decomposition of the measured wavefronts for three DM configurations, using the Noll indexing convention. In configuration A, where all DM actuators were set to 0 V, the modal energy is concentrated in the defocus mode (Noll index 5), with negligible contributions from higher-order aberrations. Configuration B corresponds to the optimized wavefront obtained through MSBO. Here, the modal energy is more broadly distributed, primarily across low-order modes such as tilt (indices 2–3), astigmatism and coma (indices 6–8), suggesting that the optimization process selectively compensated for defocus (index 4) while introducing other types of aberrations. This suggests that non-ideal, structured focal shapes — containing astigmatism, slight tilt, coma or higher-order features — can improve energy coupling to hot electrons and thereby boost the proton energy. Rather than seeking a perfectly flat wavefront, the optimization exploits constructive asym-

metry to concentrate energy into favorable regions at the target, improving laser-to-proton energy conversion. Configuration C, based on a WFS-guided correction, exhibits a lower modal energy profile compared to B with reduced amplitude in mid-order modes (indices >10). This indicates that the WFS correction effectively reduces overall aberration levels but may lack the fine-tuning capabilities of MSBO towards improving the laser energy deposition on the target.

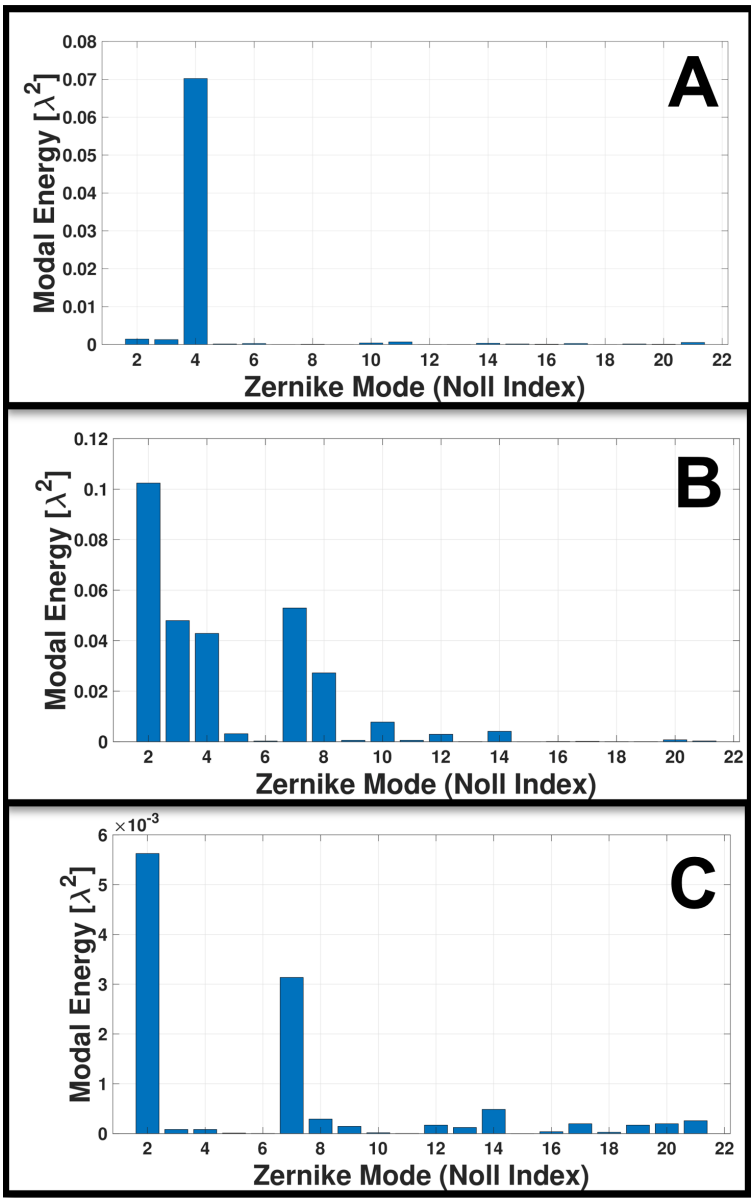

FIGURE 5.21 : Zernike modal energy three DM configurations.

Modal decomposition of the measured wavefronts into Zernike polynomials (Noll indexing) for three DM configurations. The bar plots represent the modal energy [λ^2] associated with the first 21 Zernike modes. (A) DM actuators set to zero volts. (B) MSBO DM configuration and (C) WFS-based DM configuration.

5.12 Discussion

For MSBO to be effective, source stability is crucial, as experimental noise inherently limits prediction accuracy. Indeed, the experimental standard deviation, the standard deviation of the predicted maximum energy, and the RMSE of the RF model, are all of the same order ($\sim$5%), indicating that the model's predictive accuracy is well-aligned with the shot-to-shot fluctuations.

In our case, each target-holder allows for approximately 400 shots before target-holder replacement, with the optimization process requiring $\sim$150 shots, leaving 250 for further studies. Even with only 50–100 shots, MSBO maintains good performance while preserving more shots for additional studies. Testing the MSBO-determined DM configuration on a new target-holder resulted in a lower maximum proton energy ($\mathcal{E}_{p,max} = (5.82 \pm 0.38)$ MeV), a 13% reduction, with minimal misalignment effects. This decline is likely due to target surface quality, alignment variations, or long-term laser energy drift, which can reach $\sim$10% if unadjusted during an experiment.

The superior performance of non-flat wavefronts optimized via DM configurations over traditional flat profiles in laser-driven ion acceleration suggests that unconventional beam profiles could unlock new operational regimes in laser-plasma interactions. This matches with prior findings where ion acceleration optimization was achieved through advanced control techniques without maximizing on-target laser intensity (Loughran et al., 2023). Beyond peak intensity, factors such as wavefront shaping parameter tuning and spatiotemporal profiles play a crucial role in enhancing acceleration. Recent studies favor novel laser modes, such as LG beams (Wang, 2024; Brabetz et al., 2015), which enhance proton acceleration by reducing divergence and improving energy deposition. Willim et al. (2023) investigate the use of high-intensity twisted laser beams — specifically, LG modes carrying OAM — to produce collimated proton bunches with energies up to 40 MeV and a divergence approximately 6.5 times smaller than those driven by conventional Gaussian laser beams of the same energy. In fact, LG beams enhance hot electron production, leading to higher proton energies, likely due to a structured intensity distribution improving energy coupling into plasma electrons. This structured intensity profile, similar to ours, may reduce hot electron divergence, keeping more electrons near the propagation axis and ensuring a uniform sheath field for efficient proton acceleration.

Our findings demonstrate how non-Gaussian laser modes improve energy coupling to plasma, enhancing proton beam metrics. Understanding the specific focal spot shape emerging from this optimization would require costly 2D–3D simulations to investigate its influence on energy transfer and plasma dynamics in the TNSA process. While such studies could provide deeper insights, this study primarily demonstrates MSBO's effectiveness in tuning maximum proton energy for applications and identifying key areas for future in-depth research.

While the primary goal of our MSBO algorithm was to maximize the proton beam metrics, our experimental observations do not show an improvement in energy stability when operating under the

optimized actuator configuration. Several known sources contribute to the shot-to-shot fluctuations in proton maximum energy, including :

— **Laser energy fluctuations :** The ALLS laser exhibits a 2.5% RMS energy variation, which directly influences the intensity on target and hence the accelerating electric field.

— **Long-term thermal wavefront drift :** Over time, thermal effects within the laser chain can cause slow distortions of the wavefront, altering the focal spot shape.

— **Target non-planarity :** Minor deformations or surface irregularities in the target introduce inconsistencies in laser–plasma coupling, which affect the efficiency and stability of ion acceleration.

Although the percentage fluctuation remains comparable and inherent to the experimental noise, the optimized configuration maintains shot-to-shot stability while operating in a higher-energy regime, suggesting that MSBO identifies not only high-energy configurations but also robust solutions within the experimental noise envelope.

While our optimization targeted the maximum proton energy, we also examined other available beam metrics presented in Table 5.5. Beam divergence was not measured in this experiment and therefore cannot be evaluated. Regarding N_{tot}, we found that this metric exhibits significant shot-to-shot fluctuations (on the order of $\sim$18%), making precise optimization more challenging. Nevertheless, N_{tot} is mostly improved for proton energies above 3 MeV as it can be observed in Figure 5.12.A. The MSBO configuration outperformed the others in nearly all aspects. It produced the highest number of protons in the 3–8 MeV energy range at 1.61×10^{10} protons.sr^{-1}, representing a 77% enhancement over WFS correction and a fivefold increase compared to the zeroed DM actuators configuration. In terms of energy transfer, MSBO achieved the highest laser-to-proton energy efficiency at 0.74%, versus 0.61% for WFS correction and 0.60% for zeroed DM actuators. The mean proton energy under MSBO reached 1.68 MeV, surpassing WFS correction (1.41 MeV) and zeroed DM actuators (1.23 MeV), while the second moment (indicative of spectral width) was also highest for MSBO at 4.63 MeV^2. Finally, the hot electron temperature, extracted from the logarithmic slope of the energy spectrum, reached 1.58 MeV for MSBO. This is significantly higher than the values for WFS correction (1.16 MeV) and zeroed DM actuators (0.78 MeV), suggesting more effective electron heating in the MSBO configuration. Altogether, these findings highlight the crucial role of wavefront shaping and data-driven optimization in enhancing both the energy and quality of laser-driven proton beams.

TABLE 5.5 : Proton beam metrics across three DM configurations.

Metric	MSBO	WFS	Zeroed DM actuators
N_{tot}^{3-8MeV} [protons·sr^{-1}]	1.61×10^{10} ($\pm 20\%$)	9.12×10^{9} ($\pm 17\%$)	3.12×10^{9} ($\pm 18\%$)
$\mathcal{E}_{tot}$ [J]	2.35×10^{-2} ($\pm 6\%$)	1.94×10^{-2} ($\pm 15\%$)	1.93×10^{-2} ($\pm 14\%$)
α [%]	0.74 ($\pm 6\%$)	0.61 ($\pm 15\%$)	0.60 ($\pm 14\%$)
$\bar{\mathcal{E}}$ [MeV]	1.68 ($\pm 2\%$)	1.41 ($\pm 2\%$)	1.23 ($\pm 2\%$)
M^2 [MeV^2]	4.63 ($\pm 9\%$)	3.02 ($\pm 4\%$)	2.06 ($\pm 5\%$)
T_{hot} [MeV]	1.58 ($\pm 8\%$)	1.16 ($\pm 6\%$)	0.78 ($\pm 7\%$)

5.13 Conclusion

Our approach is particularly advantageous for facilities with limited diagnostic capabilities, as it eliminates the need for complex wavefront measurements at nominal energy. By leveraging data-driven optimization it provides a cost-effective and scalable pathway for enhancing laser-driven ion acceleration, making performance gains accessible to a broader range of high-power laser installations.

In this study, we focused on optimizing 20 of the 48 DM actuators to maintain computational efficiency. Scaling to a larger number of control parameters remains challenging for conventional BO, as surrogate models tend to lose accuracy in high-dimensional spaces (Frazier, 2018). Nevertheless, advanced surrogate modeling, parallelized BO strategies, and hybrid optimization schemes offer promising routes to overcome these limitations and could ultimately unlock higher proton energies. Furthermore, adapting this methodology to facilities with low repetition rates or severe shot constraints (e.g., fewer than 100 shots) would broaden its applicability.

Overall, this framework demonstrates a resource-efficient route for accelerating optimization in laser-plasma experiments, paving the way toward more systematic and reliable tuning strategies. Its versatility opens the door to future extensions targeting other laser-plasma processes, reinforcing data-driven optimization as a key enabler for next-generation high-intensity laser science.

6 GENERAL CONCLUSION

6.1 Summary and novelty of the work

This doctoral work focused on the development, implementation, and experimental validation of Machine Learning (ML)-assisted optimization strategies for laser-driven ion acceleration in the Target Normal Sheath Acceleration (TNSA) regime, using the 150 TW Ti : Sapphire laser at the *Advanced Laser Light Source* (*ALLS*). The overall goal was to demonstrate that data-driven optimization methods, and in particular surrogate-assisted Bayesian Optimization (BO), can be used to efficiently explore and exploit high-dimensional experimental parameter spaces, even under the challenging constraints of low repetition rate (0.625 Hz), significant shot-to-shot fluctuations, and measurement noise typical of high-power laser-plasma experiments.

The work addressed three major, interconnected objectives :

1. **Establishing a robust and repeatable experimental platform.** A prerequisite for any meaningful optimization is the ability to deliver reproducible experimental conditions over many consecutive laser shots. This thesis therefore began by stabilizing the *ALLS* ion beamline through :
   - The design and deployment of a high-capacity (400 targets) automated target delivery system, enabling extended operation without manual intervention.
   - The implementation of precise target alignment routines, ensuring consistent focal positioning at the target interaction point.
   - The integration of synchronized data acquisition and beam diagnostics to monitor key beam parameters in real time.

   These developments reduced the relative shot-to-shot fluctuations in maximum proton energy and in proton yield, establishing a stable baseline from which to measure the impact of optimization strategies.
2. **Systematic investigation of spatio-temporal laser shaping effects on TNSA performance.** A combined adaptive optics and temporal shaping setup was commissioned, consisting of a bimorph deformable mirror (DM) with 20 out of 48 actuators for wavefront control, and an acousto-optic programmable dispersive filter (AOPDF) for spectral phase modulation. Using this system, the thesis quantified :
   - The relative influence of spatial (wavefront) vs. temporal (dispersion) shaping on proton beam characteristics for thin (4.5 μm) Al foils.
   - The robustness of maximum proton energy as a single-parameter metric for optimization, due to its strong correlation with other beam quality indicators (yield, spectral shape), and low shot-to-shot variability.

— The limits of conventional flat wavefront optimization approaches, showing that non-planar optimal wavefronts can outperform maximum-intensity configurations.

3. **Development and validation of a multi-step surrogate-assisted Bayesian Optimization (MSBO) framework.** The core methodological contribution of the thesis is the implementation of a modular optimization framework leveraging ML techniques that :
   — Uses initial experimental data to train a surrogate model (Random Forest regression) predicting maximum proton energy from actuator voltage configurations.
   — Applies an MSBO strategy to efficiently search the reduced parameter space (20 selected actuators out of 48 for computational efficiency).
   — Iteratively updates the surrogate model with experimental feedback, refining predictions and guiding subsequent actuator configurations.

   This approach achieved :
   — A 72% gain in maximum proton energy compared to the zeroed DM actuators' configuration.
   — A 24% improvement over standard wavefront sensor-based correction methods.
   — Convergence within fewer than 150 laser shots : a critical requirement for low-repetition, high-cost experiments.

   Beyond the achieved performance gains, the MSBO results revealed that the optimal DM configuration introduced structured aberrations rather than removing them entirely, providing interesting physical insights into the complex laser-plasma coupling dynamics in TNSA.

In summary, this thesis advances the state-of-the-art in laser-driven ion acceleration by :

— Delivering a stable, automated, and high-capacity experimental platform for optimization experiments and multidisciplinary applications.
— Providing quantitative spatio-temporal sensitivity mapping of TNSA performance using wavefront and spectral control.
— Demonstrating that surrogate-assisted BO can yield performance gains with minimal shot counts, while uncovering non-intuitive optimal configurations.

6.2 Perspectives and future directions

The work presented in this thesis establishes a robust experimental and methodological foundation for the integration of ML-assisted optimization in laser-driven ion acceleration. The approaches developed are inherently platform-agnostic and can be extended to other high-intensity laser–matter interaction regimes, such as LWFA, and laser-driven neutron, gamma-ray, or X-ray sources. Building on this foundation, several avenues emerge for advancing the capabilities toward autonomous, adaptive, and multi-regime operation :

1. **Transition from offline to online optimization.** The optimization strategies presented here were primarily implemented in an offline configuration, relying on surrogate models trained from pre-acquired datasets. A natural progression is the integration of the BO framework into a real-time, closed-loop environment. Coupling live diagnostic feedback with adaptive hyperparameter tuning would allow the system to account for experimental drifts, laser pointing fluctuations, and shot-to-shot variability. This step is expected to enhance convergence speed, resilience to noise, and reproducibility of beam performance directly *in situ*.
2. **Parametric exploration and feature selection.** Applying this methodology to other platforms requires a systematic identification of the parameters most critical to performance. Techniques such as correlation analysis, variance-based sensitivity indices, and advanced feature selection algorithms can be employed to reduce dimensionality and improve interpretability. This will guide experimental prioritization, ensuring that optimization efforts are concentrated on the degrees of freedom with the highest influence on target metrics.
3. **Optimization of advanced radiation sources.** Beyond ion acceleration, the framework can be generalized to optimize LWFA-based electron beams, betatron X-rays, and secondary particle sources such as neutrons or gamma rays. Both offline and online optimization modes could be used to enhance stability, reproducibility, and adaptability across high-dimensional parameter spaces, even under conditions of limited diagnostic coverage. The long-term goal is to deliver application-ready, high-quality radiation sources across multiple regimes.
4. **Real-time adaptive control systems.** Developing a modular software–hardware architecture capable of monitoring key performance indicators — such as beam pointing, stability, and spectral properties — in real time and applying corrective actions autonomously represents a transformative step. Approaches based on reinforcement learning (RL) could enable continuous optimization during extended experimental campaigns, rapid adaptation to evolving laser or target conditions, and ultimately, fully autonomous operation.

In the longer term, these advances could transform laser–plasma accelerators into self-optimizing systems, drastically shortening tuning procedures, boosting experimental efficiency, and delivering high-quality beams for a wide spectrum of applications. By lowering operational complexity while improving beam stability and performance, this approach lays the groundwork for a new generation of compact, reliable, and versatile sources, bridging the gap between proof-of-principle experiments and practical implementation in many fields.

6.3 Concluding remarks

This thesis establishes ML-guided optimization as a practical and effective route to improving TNSA performance under realistic experimental constraints. Beyond the demonstrated gain in TNSA acceleration with limited shots, the work provides a scalable framework for adaptive, data-driven control of high-intensity laser-plasma interactions. By extending to online control, adding complementary optimization methods, and generalizing across regimes, this approach charts a credible

path toward autonomous, application-ready laser-plasma sources for science, technology, and society.

BIBLIOGRAPHY

Afshari M, Hornung J, Kleinschmidt A, Neumayer P, Bertini D, Bagnoud V (2020) Proton acceleration via the TNSA mechanism using a smoothed laser focus. *AIP Advances*, 10(3). https://pubs.aip.org/aip/adv/article/10/3/035023/103688800.

Agosteo S, Anania M, Caresana M, Cirrone G, De Martinis C, Delle Side D, Fazzi A, Gatti G, Giove D, Giulietti D et al. (2014) The LILIA (laser induced light ions acceleration) experiment at LNF. *Nuclear Instruments and Methods in Physics Research Section B : Beam Interactions with Materials and Atoms*, 331:15–19. https://www.sciencedirect.com/science/article/pii/S0168583X14001207.

Albert F, Thomas A, Mangles S, Banerjee S, Corde S, Flacco A, Litos M, Neely D, Vieira J, Najmudin Z et al. (2014) Laser wakefield accelerator based light sources : potential applications and requirements. *Plasma Physics and Controlled Fusion*, 56(8):084015. https://iopscience.iop.org/article/10.1088/0741-3335/56/8/084015/meta.

Albert F, Thomas AG (2016) Applications of laser wakefield accelerator-based light sources. *Plasma Physics and Controlled Fusion*, 58(10):103001. https://iopscience.iop.org/article/10.1088/0741-3335/58/10/103001/meta.

Albertazzi B, Ciardi A, Nakatsutsumi M, Vinci T, Béard J, Bonito R, Billette J, Borghesi M, Burkley Z, Chen S, Cowan T, Herrmannsdörfer T, Higginson D, Kroll F, Pikuz S, Naughton K, Romagnani L, Riconda C, Revet G, Fuchs J (2014) Laboratory formation of a scaled protostellar jet by coaligned poloidal magnetic field. *Science (New York, N.Y.)*, 346:325–8. https://www.science.org/doi/full/10.1126/science.1259694.

Alejo A, Kar S, Tebartz A, Ahmed H, Astbury S, Carroll D, Ding J, Doria D, Higginson A, McKenna P et al. (2016) High resolution Thomson Parabola Spectrometer for full spectral capture of multi-species ion beams. *Review of Scientific Instruments*, 87(8). https://pubs.aip.org/aip/rsi/article/87/8/083304/839186.

Aleksandrov A, Zavalova V, Kudryashov A, Rukosuev A, Romanov P, Samarkin V (2004) Closed adaptive systems with controllable bimorph mirrors. *Journal of optical technology*, 71(11):737–741. https://opg.optica.org/ao/fulltext.cfm?uri=ao-27-7-1223&id=165969.

Amiranoff F, Baton S, Bernard D, Cros B, Descamps D, Dorchies F, Jacquet F, Malka V, Marques J, Matthieussent G et al. (1998) Observation of laser wakefield acceleration of electrons. *Physical Review Letters*, 81(5):995. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.81.995.

Amodio A, Wang D, Berger C, Tsai HE, Barber SK, van Tilborg J, Picksley A, Eisentraut Z, Vora NR, Logantha M et al. (2025) Pointing stabilization of a 1 Hz high-power laser via machine learning. *High Power Laser Science and Engineering*, 13:e35. https://www.cambridge.org/core/journals/high-power-laser-science-and-engineering/article/pointing-stabilization-of-a-1hz-high-power-laser-via-machine-learning/6F9BD740921C6A445967A8F42FD7A7CB.

Amplitude Laser (2019) *DAZZLER^{TM} system manual Part 1 : installation operation*. https://www.wavequanta.com/product/desc/id/220.html. https://www.wavequanta.com/product/desc/id/220.html.

Amplitude Laser (2025a) *Sequoia — Instrumentation Femtosecond Laser (third-order cross correlator)*. https://amplitude-laser.com/products/femtosecond-lasers/instrumentation-lasers-femtosecondes/sequoia/. Accessed : 2025-10-16.

Amplitude Laser (2025b) *Wizzler – Single-shot spectral phase measurement device*. https://amplitude-laser.com/fr/produits/lasers-femtosecondes/instrumentation-fr/wizzler/. Accessed : July 30, 2025.

Andreev N, Gorbunov L, Kirsanov V, Nakajima K, Ogata A (1997) Structure of the wake field in plasma channels. *Physics of Plasmas*, 4(4):1145–1153. https://pubs.aip.org/aip/pop/article/4/4/1145/263475/Structure-of-the-wake-field-in-plasma-channels.

Anirudh R, Archibald R, Asif MS, Becker MM, Benkadda S, Bremer PT, Budé RH, Chang CS, Chen L, Churchill R et al. (2023) 2022 review of data-driven plasma science. *IEEE Transactions on Plasma Science*, 51(7):1750–1838. https://ieeexplore.ieee.org/abstract/document/10214236.

Antici P, Fazi M, Lombardi A, Migliorati M, Palumbo L, Audebert P, Fuchs J (2008) Numerical study of a linear accelerator using laser-generated proton beams as a source. *Journal of Applied Physics*, 104(12):124901. https://doi.org/10.1063/1.3021160.

Antici P, Migliorati M, Mostacci A, Picardi L, Palumbo L, Ronsivalle C (2011) A compact post-acceleration scheme for laser-generated protons. *Physics of Plasmas*, 18(7):073103. https://doi.org/10.1063/1.3574361.

Awad M, Khanna R, Awad M, Khanna R (2015) Support vector regression. *Efficient learning machines : Theories, concepts, and applications for engineers and system designers*, pages 67–80. https://link.springer.com/chapter/10.1007/978-1-4302-5990-9_4.

Baduge SK, Thilakarathna S, Perera JS, Arashpour M, Sharafi P, Teodosio B, Shringi A, Mendis P (2022) Artificial intelligence and smart vision for building and construction 4.0 : Machine and deep learning methods and applications. *Automation in Construction*, 141:104440. https://ars.els-cdn.com/content/image/1-s2.0-S0926580522003132-ga1_lrg.jpg.

Baer CRE, Kränkel C, Saraceno CJ, Heckl OH, Golling M, Peters R, Petermann K, Südmeyer T, Huber G, Keller U (2010) Femtosecond thin-disk laser with 141 W of average power. *Optics letters*, 35(13):2302–2304. https://opg.optica.org/ol/fulltext.cfm?uri=ol-35-13-2302&id=203444.

Baraniuk C (2023) *The World's Most Powerful Lasers*. https://www.bbc.com/future/article/20231123-the-worlds-most-powerful-lasers. Accessed : 2025-03-31.

Barberio M, Antici P (2019) Laser-PIXE using laser-accelerated proton beams. *Scientific Reports*, 9:6855. https://www.nature.com/articles/s41598-019-42997-y.

Barberio M, Giusepponi S, Vallières S, Scisció M, Celino M, Antici P (2020) Ultra-fast high-precision metallic nanoparticle synthesis using laser-accelerated protons. *Scientific Reports*, 10(1):9570. https://www.nature.com/articles/s41598-020-65282-9.

Barberio M, Scisciò M, Vallières S, Cardelli F, Chen S, Famulari G, Gangolf T, Revet G, Schiavi A, Senzacqua M et al. (2018a) Laser-accelerated particle beams for stress testing of materials. *Nature communications*, 9(1):372. https://www.nature.com/articles/s41467-017-02675-x.

Barberio M, Scisciò M, Vallières S, Veltri S, Morabito A, Antici P (2017a) Laser-generated proton beams for high-precision ultra-fast crystal synthesis. *Scientific Reports*, 7(1):12522. https://www.nature.com/articles/s41598-017-12782-w.

Barberio M, Sciscio M, Veltri S, Antici P (2016) Fabrication of nanostructured targets for improved laser-driven proton acceleration. *Superlattices and Microstructures*, 95:159–163. https://www.sciencedirect.com/science/article/pii/S0749603616301653.

Barberio M, Vallières S, Scisciò M, Kolhatkar G, Ruediger A, Antici P (2018b) Graphitization of diamond by laser-accelerated proton beams. *Carbon*, 139:531–537. https://www.sciencedirect.com/science/article/pii/S0008622318305876.

Barberio M, Veltri S, Sciscio M, Antici P (2017b) Laser-accelerated proton beams as diagnostics for cultural heritage. *Scientific Reports*, 7(1):40415. https://www.nature.com/articles/s41598-017-12782-w.

Barty C, Gordon Iii C, Lemoff B (1994) Multiterawatt 30-fs Ti : sapphire laser system. *Optics letters*, 19(18):1442–1444. https://opg.optica.org/ol/fulltext.cfm?uri=ol-19-18-1442&id=12561.

Barty C, Key M, Britten J, Beach R, Beer G, Brown C, Bryan S, Caird J, Carlson T, Crane J et al. (2004) An overview of LLNL high-energy short-pulse technology for advanced radiography of laser fusion experiments. *Nuclear Fusion*, 44(12):S266. https://iopscience.iop.org/article/10.1088/0029-5515/44/12/S18/meta.

Basak D, Pal S, Patranabis DC et al. (2007) Support vector regression. *Neural Information Processing-Letters and Reviews*, 11(10):203–224. https://www.researchgate.net/profile/Srimanta-Pal/publication/228537532_Support_Vector_Regression/links/00b7d51c2a0d86d3d9000000/Support-Vector-Regression.pdf.

Beaurepaire B, Vernier A, Bocoum M, Böhle F, Jullien A, Rousseau J, Lefrou T, Douillet D, Iaquaniello G, Lopez-Martens R et al. (2015) Effect of the laser wavefront in a laser-plasma accelerator. *Physical Review X*, 5(3):031012. https://journals.aps.org/prx/abstract/10.1103/PhysRevX.5.031012.

Beier N, Senthilkumaran V, Kriz E, Fourmaux S, Légaré F, Ma T, Hussein A (2024) Deep learning based x-ray spectrometer for high repetition rate characterization of betatron radiation. *Physics of Plasmas*, 31(10). https://pubs.aip.org/aip/pop/article/31/10/103106/3315998.

Biasi R, Gallieni D, Salinari P, Riccardi A, Mantegazza P (2010) Contactless thin adaptive mirror technology : past, present, and future. *Adaptive Optics Systems II*, SPIE, volume 7736, pages 872–885. https://doi.org/10.1117/12.858816.

Bifano T (2011) MEMS deformable mirrors. *Nature photonics*, 5(1):21–23. https://www.nature.com/articles/nphoton.2010.297#citeas.

Bisesto F, Galletti M, Anania MP, Ferrario M, Pompili R, Botton M, Schleifer E, Zigler A (2019) Review on TNSA diagnostics and recent developments at SPARC_LAB. *High Power Laser Science and Engineering*, 7:e56. https://www.cambridge.org/core/journals/high-power-laser-science-and-engineering/article/review-on-tnsa-diagnostics-and-recent-developments-at-sparclab/B440C51C869AED1D23AC766AF28C8AAA.

Bishop CM, Nasrabadi NM (2006) *Pattern recognition and machine learning*. volume 4. Springer. https://link.springer.com/book/9780387310732.

Boivin F, Vallières S, Fourmaux S, Payeur S, Antici P (2022) Quantitative laser-based x-ray fluorescence and particle-induced x-ray emission. *New Journal of Physics*, 24(5):053018. https://dx.doi.org/10.1088/1367-2630/ac6767.

Bolton PR, Parodi K, Schreiber J (2018) Overview of Applications of Laser-Driven Particle Acceleration (Editors Paul R. Bolton, Katia Parodi, and Jörg Schreiber) by CRC Press (Taylor and Francis Group) ISBN 9781498766418—5 June 2018. *Quantum Beam Science*, 2(4):25. https://www.mdpi.com/2412-382X/2/4/25.

Borghesi M, Bigongiari A, Kar S, Macchi A, Romagnani L, Audebert P, Fuchs J, Toncian T, Willi O, Bulanov SV, Mackinnon AJ, Gauthier JC (2008) Laser-driven proton acceleration : source optimization and radiographic applications. *Plasma Physics and Controlled Fusion*, 50(12): 124040. https://dx.doi.org/10.1088/0741-3335/50/12/124040.

Borghesi M, Campbell D, Schiavi A, Haines M, Willi O, MacKinnon A, Patel P, Gizzi L, Galimberti M, Clarke R et al. (2002) Electric field detection in laser-plasma interaction experiments via the proton imaging technique. *Physics of Plasmas*, 9(5):2214–2220. https://pubs.aip.org/aip/pop/article/9/5/2214/1070312/Electric-field-detection-in-laser-plasma.

Borghesi M, Fuchs J, Bulanov S, Mackinnon A, Patel P, Roth M (2006) Fast ion generation by high-intensity laser irradiation of solid targets and applications. *Fusion science and technology*, 49(3):412–439. https://www.tandfonline.com/doi/abs/10.13182/FST06-A1159.

Borghesi M, Mackinnon A, Campbell DH, Hicks D, Kar S, Patel PK, Price D, Romagnani L, Schiavi A, Willi O (2004) Multi-MeV proton source investigations in ultraintense laser-foil interactions. *Physical Review Letters*, 92(5):055003. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.92.055003.

Brabetz C, Busold S, Cowan T, Deppert O, Jahn D, Kester O, Roth M, Schumacher D, Bagnoud V (2015) Laser-driven ion acceleration with hollow laser beams. *Physics of Plasmas*, 22(1). https://pubs.aip.org/aip/pop/article-abstract/22/1/013105/318240/Laser-driven-ion-acceleration-with-hollow-laser?redirectedFrom=fulltext.

Breiman L (2001) Random forests. *Machine learning*, 45(1):5–32. https://link.springer.com/article/10.1023/a:1010933404324.

Brenner C, Green J, Robinson A, Carroll D, Dromey B, Foster P, Kar S, Li Y, Markey K, Spindloe C, Streeter M, Tolley M, Wahlström CG, Xu M, Zepf M, McKenna P, Neely D (2011) Dependence of laser accelerated protons on laser energy following the interaction of defocused, intense laser pulses with ultra-thin targets. *Laser and Particle Beams*, 29:345–351. https://pubs.aip.org/aip/apl/article/104/8/081123/132539.

Brida D, Manzoni C, Cerullo G (2007) Phase-locked pulses for two-dimensional spectroscopy by a birefringent delay line. *Optics Letters*, 32(20):2891–2893. https://opg.optica.org/ol/fulltext.cfm?uri=ol-37-15-3027&id=239856.

Brochu E, Cora VM, De Freitas N (2010) A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv :1012.2599*. https://doi.org/10.48550/arXiv.1012.2599.

Bruchon N, Fenu G, Gaio G, Lonza M, O'Shea FH, Pellegrino FA, Salvato E (2020) Basic reinforcement learning techniques to control the intensity of a seeded free-electron laser. *Electronics*, 9(5):781. https://www.mdpi.com/2079-9292/9/5/781.

Brunel F (1987) Not-so-resonant, resonant absorption. *Physical review letters*, 59(1):52. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.59.52.

Bulanov S, Khoroshkov V (2002) Feasibility of using laser ion accelerators in proton therapy. *Plasma Physics Reports*, 28:453–456. https://link.springer.com/article/10.1134/1.1478534.

Cain MG, Stewart M (2014) Standards for piezoelectric and ferroelectric ceramics. *Characterisation of ferroelectric bulk materials and thin films*, Springer, pages 267–275. https://doi.org/10.1007/978-1-4020-9311-1_12.

Cakir A, Guzel O (2020) *A Brief Review of Plasma Wakefield Acceleration*. https://arxiv.org/abs/1908.07207.

Capuano F, Peceli D, Tiboni G, Camoriano R, Rus B (2023) TempoRL : laser pulse temporal shape optimization with deep reinforcement learning. *High-power, High-energy Lasers and Ultrafast Optical Technologies*, SPIE, volume 12577, pages 62–74. https://www.spiedigitallibrary.org/conference-proceedings-of-spie/12577/125770C/TempoRL--laser-pulse-temporal-shape-optimization-with-deep-reinforcement/10.1117/12.2669267.short.

Capuano F, Peceli D, Tiboni G (2025) Shaping Laser Pulses with Reinforcement Learning. *arXiv preprint arXiv :2503.00499*. https://arxiv.org/abs/2503.00499.

Catrix E, Boivin F, Langlois K, Vallières S, Boynukara CY, Fourmaux S, Antici P (2023) Stable high repetition-rate laser-driven proton beam production for multidisciplinary applications on the advanced laser light source ion beamline. *Review of Scientific Instruments*, 94(10). https://doi.org/10.1063/5.0160783.

Catrix E, Fourmaux S, Vallières S, Bianchi F, Fillion-Gourdeau F, Maltais J, MacLean S, Antici P (2025) 20-dimensional surrogate-assisted Bayesian optimization of laser-driven proton beams. *Applied Physics Letters*, 126(25):254104. https://doi.org/10.1063/5.0272789.

Ceccotti T, Lévy A, Popescu H, Réau F, d'Oliveira P, Monot P, Geindre J, Lefebvre E, Martin P (2007) Proton acceleration with high-intensity ultrahigh-contrast laser pulses. *Physical review letters*, 99(18):185002. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.99.185002.

Chagovets T, Stanček S, Giuffrida L, Velyhan A, Tryus M, Grepl F, Istokskaia V, Kantarelou V, Wiste T, Hernandez Martin JC, Schillaci F, Margarone D (2021) Automation of Target Delivery and Diagnostic Systems for High Repetition Rate Laser-Plasma Acceleration. *Applied Sciences*, 11(4). https://www.mdpi.com/2076-3417/11/4/1680.

Chanteloup JC (2005) Multiple-wave lateral shearing interferometry for wave-front sensing. *Applied optics*, 44(9):1559–1571. https://opg.optica.org/ao/fulltext.cfm?uri=ao-44-9-1559&id=82949.

Chen S, Vranic M, Gangolf T, Boella E, Antici P, Bailly-Grandvaux M, Loiseau P, Pépin H, Revet G, Santos J, Schroer AM, Starodubtsev M, Willi O, Silva L, D'Humières E, Fuchs J (2017)

Collimated protons accelerated from an overdense gas jet irradiated by a 1 micron wavelength high-intensity short-pulse laser. *Scientific Reports*, 7. https://www.nature.com/articles/s41598-017-12910-6.

Chen SN, Gauthier M, Higginson DP, Dorard S, Mangia F, Riquier R, Atzeni S, Marquès JR, Fuchs J (2014) Monochromatic short pulse laser produced ion beam using a compact passive magnetic device. *Review of Scientific Instruments*, 85(4):043504. https://doi.org/10.1063/1.4870250.

Choobini A, Ghaffari-Oskooei S (2024) Investigation of self-focusing of Gaussian laser beams within magnetized plasma via source-dependent expansion method. *arXiv preprint arXiv :2404.00331*. https://pubs.aip.org/aip/pop/article/31/8/082105/3306661.

Cirrone GP, Cuttone G, Romano F, Schillaci F, Scuderi V, Amato A, Candiano G, Costa M, Gallo G, Larosa G et al. (2015) Design and status of the ELIMED beam line for laser-driven ion beams. *Applied Sciences*, 5(3):427–445. https://www.mdpi.com/2076-3417/5/3/427.

Clark E, Krushelnick K, Davies J, Zepf M, Tatarakis M, Beg F, Machacek A, Norreys P, Santala M, Watts I et al. (2000) Measurements of energetic proton transport through magnetized plasma from intense laser interactions with solids. *Physical Review Letters*, 84(4):670. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.84.670.

Clayton C, Marsh K, Dyson A, Everett M, Lal A, Leemans W, Williams R, Joshi C (1993) Ultrahigh-gradient acceleration of injected electrons by laser-excited relativistic electron plasma waves. *Physical review letters*, 70(1):37. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.70.37.

Clayton CE, Ralph J, Albert F, Fonseca R, Glenzer S, Joshi C, Lu W, Marsh K, Martins SF, Mori WB et al. (2010) Self-guided laser wakefield acceleration beyond 1 GeV using ionization-induced injection. *Physical review letters*, 105(10):105003. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.105.105003.

Cobo C, Arran C, Bourgeois N, Calvin L, Carderelli J, Cavanagh N, Colgan C, Dann S, Fitzgarrald R, Gerstmayr E et al. (2024) Target sensitivity study of density transition-injected electrons in laser wakefield accelerators. *Physical Review Accelerators and Beams*, 27(11):111301. https://journals.aps.org/prab/abstract/10.1103/PhysRevAccelBeams.27.111301.

Contributors SO (2025) *skopt.gp_minimize — Scikit-Optimize Documentation*. https://scikit-optimize.github.io/stable/modules/generated/skopt.gp_minimize.html. Accessed : 2025-01-17.

Copplestone SM, Pfeiffer M, Fasoulas S, Munz CD (2019) High-order Particle-In-Cell simulations of laser-plasma interaction. *The European Physical Journal Special Topics*, 227:1603–1614. https://link.springer.com/article/10.1140/epjst/e2019-800160-y.

Corde S, Ta Phuoc K, Lambert G, Fitour R, Malka V, Rousse A, Beck A, Lefebvre E (2013) Femtosecond x rays from laser-plasma accelerators. *Reviews of Modern Physics*, 85(1):1–48. https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.85.1.

Coury M, Carroll D, Robinson A, Yuan X, Brenner C, Burza M, Gray R, Lancaster K, Li Y, Lin X et al. (2013) Injection and transport properties of fast electrons in ultraintense laser-solid interactions. *Physics of Plasmas*, 20(4). https://pubs.aip.org/aip/pop/article/20/4/043104/1032465.

Coury M, Carroll DC, Robinson APL, Yuan XH, Brenner CM, Burza M, Gray RJ, Quinn MN, Lancaster KL, Li YT, Lin XX, Tresca O, Wahlström CG, Neely D, McKenna P (2012) Influence of laser irradiated spot size on energetic electron injection and proton acceleration in foil targets. *Applied Physics Letters*, 100(7):074105. https://doi.org/10.1063/1.3685615.

Cowan T, Fuchs J, Ruhl H, Kemp A, Audebert P, Roth M, Stephens R, Barton I, Blazevic A, Brambrink E et al. (2004) Ultralow emittance, multi-MeV proton beams from a laser virtual-cathode plasma accelerator. *Physical review letters*, 92(20):204801. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.92.204801.

Cugat O, Basrour S, Divoux C, Mounaix P, Reyne G (2001) Deformable magnetic mirror for adaptive optics : technological aspects. *Sensors and Actuators A : Physical*, 89(1-2):1–9. https://www.sciencedirect.com/science/article/pii/S0924424700005501.

Dabu R (2018) High-Power, High-Intensity Contrast Hybrid Femtosecond Laser Systems. *High Power Laser Systems*, 43 pages. https://www.intechopen.com/chapters/56920.

Daido H, Nishiuchi M, Pirozhkov AS (2012) Review of laser-driven ion sources and their applications. *Reports on progress in physics*, 75(5):056401. https://iopscience.iop.org/article/10.1088/0034-4885/75/5/056401/meta.

Danson CN, Haefner C, Bromage J, Butcher T, Chanteloup JCF, Chowdhury EA, Galvanauskas A, Gizzi LA, Hein J, Hillier DI et al. (2019) Petawatt and exawatt class lasers worldwide. *High Power Laser Science and Engineering*, 7:e54. https://www.cambridge.org/core/journals/high-power-laser-science-and-engineering/article/petawatt-and-exawatt-class-lasers-worldwide/D85E3F002CBFC286C08076E127BB5F5C.

Desai R, Zhang T, Felice JJ, Oropeza R, Smith JR, Kryshchenko A, Orban C, Dexter ML, Patnaik AK (2024) Applying Machine-Learning Methods to Laser Acceleration of Protons : Lessons Learned From Synthetic Data. *Contributions to Plasma Physics*, e202400080 pages. https://onlinelibrary.wiley.com/doi/full/10.1002/ctpp.202400080.

Dietz T, Jenne M, Bauer D, Scharun M, Sutter D, Killi A (2020) Ultrafast thin-disk multi-pass amplifier system providing 1.9 kW of average output power and pulse energies in the 10 mJ range at 1 ps of pulse duration for glass-cleaving applications. *Optics Express*, 28(8):11415–11423. https://opg.optica.org/oe/fulltext.cfm?uri=oe-28-8-11415&id=429701.

Djordjević B, Kemp A, Kim J, Ludwig J, Simpson R, Wilks S, Ma T, Mariscal D (2021a) Characterizing the acceleration time of laser-driven ion acceleration with data-informed neural networks. *Plasma Physics and Controlled Fusion*, 63(9):094005. https://iopscience.iop.org/article/10.1088/1361-6587/ac172a/meta.

Djordjević B, Kemp A, Kim J, Simpson R, Wilks S, Ma T, Mariscal D (2021b) Modeling laser-driven ion acceleration with deep learning. *Physics of Plasmas*, 28(4). https://pubs.aip.org/aip/pop/article/28/4/043105/263490.

Doherty A, Fourmaux S, Astolfo A, Ziesche R, Wood J, Finlay O, Stolp W, Batey D, Manke I, Légaré F et al. (2023) Femtosecond multimodal imaging with a laser-driven X-ray source. *Communications Physics*, 6(1):288. https://www.nature.com/articles/s42005-023-01412-9.

Dolier E, King M, Wilson R, Gray R, McKenna P (2022) Multi-parameter Bayesian optimisation of laser-driven ion acceleration in particle-in-cell simulations. *New Journal of Physics*, 24(7):073025. https://iopscience.iop.org/article/10.1088/1367-2630/ac7db4/meta.

Dollar F, Zulick C, Matsuoka T, McGuffey C, Bulanov S, Chvykov V, Davis J, Kalinchenko G, Petrov G, Willingale L et al. (2013) High contrast ion acceleration at intensities exceeding 1021 Wcm- 2. *Physics of Plasmas*, 20(5). https://pubs.aip.org/aip/pop/article/20/5/056703/109589.

Dong Q, Sheng ZM, Yu M, Zhang J (2003) Optimization of ion acceleration in the interaction of intense femtosecond laser pulses with ultrathin foils. *Physical Review E*, 68(2):026408. https://journals.aps.org/pre/abstract/10.1103/PhysRevE.68.026408.

Döpp A, Eberle C, Howard S, Irshad F, Lin J, Streeter M (2023) Data-driven science and machine learning methods in laser–plasma physics. *High Power Laser Science and Engineering*, 11:e55. https://www.cambridge.org/core/journals/high-power-laser-science-and-engineering/article/datadriven-science-and-machine-learning-methods-in-laserplasma-physics/B50C69868941B26062ECF6AFCF2BF3B9.

Dorrer C, Joffre M (2001) Characterization of the spectral phase of ultrashort light pulses. *Comptes Rendus de l'Académie des Sciences-Series IV-Physics*, 2(10):1415–1426. https://www.sciencedirect.com/science/article/abs/pii/S1296214701012793.

Dover N, Nishiuchi M, Sakaki H, Kondo K, Alkhimova M, Faenov AY, Hata M, Iwata N, Kiriyama H, Koga J et al. (2020a) Effect of small focus on electron heating and proton acceleration in ultrarelativistic laser-solid interactions. *Physical review letters*, 124(8):084802. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.124.084802.

Dover N, Nishiuchi M, Sakaki H, Kondo K, Lowe H, Alkhimova M, Ditter E, Ettlinger O, Faenov AY, Hata M et al. (2020b) Demonstration of repetitive energetic proton generation by ultra-intense laser interaction with a tape target. *High Energy Density Physics*, 37:100847. https://www.sciencedirect.com/science/article/pii/S1574181820300884.

Dover NP, Ziegler T, Assenbaum S, Bernert C, Bock S, Brack FE, Cowan TE, Ditter EJ, Garten M, Gaus L et al. (2023) Enhanced ion acceleration from transparency-driven foils demonstrated at two ultraintense laser facilities. *Light : Science & Applications*, 12(1):71. https://www.nature.com/articles/s41377-023-01083-9.

Dubietis A, Jonusauskas G, Piskarskas A (1992) Powerful femtosecond pulse generation by chirped and stretched pulse parametric amplification in BBO crystal. *Optics Communications*, 88(4–6):437–440. https://www.sciencedirect.com/science/article/abs/pii/0030401892900708.

Ducret JE, Batani D, Boutoux G, Chancé A, Gastineau B, Guillard JC, Harrault F, Jakubowska K, Lantuejoul-Thfoin I, Leboeuf D et al. (2018) Calibration of the low-energy channel Thomson parabola of the LMJ-PETAL diagnostic SEPAGE with protons and carbon ions. *Review of Scientific Instruments*, 89(2). https://pubs.aip.org/aip/rsi/article/89/2/023304/362856/Calibration-of-the-low-energy-channel-Thomson.

Duris J, Kennedy D, Hanuka A, Shtalenkova J, Edelen A, Baxevanis P, Egger A, Cope T, McIntire M, Ermon S et al. (2020) Bayesian optimization of a free-electron laser. *Physical review letters*, 124(12):124801. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.124.124801.

Ehret M, de Luis D, Apiñaniz JI, Henares JL, Lera R, Pérez-Hernández JA, Puyuelo-Valdes P, Volpe L, Gatti G (2023) High-repetition-rate solid tape target delivery system for ultra-intense

laser-matter interaction at CLPU. *arXiv preprint arXiv :2302.04769*. https://arxiv.org/abs/2302.04769.

Ehret M (2016) TNSA Proton Beam Guidance with Strong Magnetic Fields Generated by Coil Targets. https://www.researchgate.net/publication/299489283_TNSA-Proton_Beam_Guidance_with_Strong_Magnetic_Fields_Generated_by_Coil_Targets.

ELI - Extreme Light Infrastructure (2011) *White book*. https://www.academia.edu/23528627/WHITEBOOK_ELI_Extreme_Light_Infrastructure_Science_and_Technology_with_Ultra_Intense_Lasers.

Ellis EM (1999) Low-cost bimorph mirrors in adaptive optics. *University of London*. https://api.semanticscholar.org/CorpusID:116853320.

Esarey E, Schroeder CB, Leemans WP (2009) Physics of laser-driven plasma-based electron accelerators. *Reviews of modern physics*, 81(3):1229–1285. https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.81.1229.

Esarey E, Shadwick B, Catravas P, Leemans W (2002) Synchrotron radiation from electron beams in plasma-focusing channels. *Physical Review E*, 65(5):056505. https://journals.aps.org/pre/abstract/10.1103/PhysRevE.65.056505.

Esirkepov T, Borghesi M, Bulanov S, Mourou G, Tajima T (2004) Highly efficient relativistic-ion generation in the laser-piston regime. *Physical review letters*, 92(17):175003. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.92.175003.

Esirkepov T, Yamagiwa M, Tajima T (2006) Laser ion-acceleration scaling laws seen in multiparametric particle-in-cell simulations. *Physical review letters*, 96(10):105001. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.96.105001.

Esirkepov TZ, Koga JK, Sunahara A, Morita T, Nishikino M, Kageyama K, Nagatomo H, Nishihara K, Sagisaka A, Kotaki H et al. (2014) Prepulse and amplified spontaneous emission effects on the interaction of a petawatt class laser with thin solid targets. *Nuclear Instruments and Methods in Physics Research Section A : Accelerators, Spectrometers, Detectors and Associated Equipment*, 745:150–163. https://www.sciencedirect.com/science/article/pii/S0168900214001077.

Estabrook K, Kruer WL (1978) Properties of resonantly heated electron distributions. *Physical Review Letters*, 40(1):42. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.40.42.

Everett M, Lal A, Gordon D, Clayton C, Marsh K, Joshi C (1994) Trapped electron acceleration by a laser-driven relativistic plasma wave. *Nature*, 368(6471):527–529. https://www.nature.com/articles/368527a0.

Faraji F, Reza M (2025) Machine learning applications to computational plasma physics and reduced-order plasma modeling : a perspective. *Journal of Physics D : Applied Physics*, 58(10):102002. https://iopscience.iop.org/article/10.1088/1361-6463/ada167/meta.

Faure J, Glinec Y, Pukhov A, Kiselev S, Gordienko S, Lefebvre E, Rousseau JP, Burgy F, Malka V (2004) A laser-plasma accelerator producing monoenergetic electron beams. *Nature*, 431(7008):541–544. https://www.nature.com/articles/nature02963.

Faure J, Rechatin C, Norlin A, Lifschitz A, Glinec Y, Malka V (2006) Controlled injection and acceleration of electrons in plasma wakefields by colliding laser pulses. *Nature*, 444(7120):737–739. https://www.nature.com/articles/nature05393.

Feister S, Nees J, Morrison J, Frische K, Orban C, Chowdhury E, Roquemore W (2014) A novel femtosecond-gated, high-resolution, frequency-shifted shearing interferometry technique for probing pre-plasma expansion in ultra-intense laser experiments. *Review of Scientific Instruments*, 85(11). https://pubs.aip.org/aip/rsi/article/85/11/11D602/356500.

Fernández EJ, Vabre L, Hermann B, Unterhuber A, Považay B, Drexler W (2006) Adaptive optics with a magnetic deformable mirror : applications in the human eye. *Optics Express*, 14(20):8900–8917. https://opg.optica.org/oe/fulltext.cfm?uri=oe-14-20-8900&id=114566.

Fernández JC, Cort Gautier D, Huang C, Palaniyappan S, Albright BJ, Bang W, Dyer G, Favalli A, Hunter JF, Mendez J et al. (2017) Laser-plasmas in the relativistic-transparency regime : Science and applications. *Physics of plasmas*, 24(5). https://pubs.aip.org/aip/pop/article/24/5/056702/992803.

Ferran Pousa A, Jalas S, Kirchen M, Martinez de la Ossa A, Thévenet M, Hudson S, Larson J, Huebl A, Vay JL, Lehe R (2023) Bayesian optimization of laser-plasma accelerators assisted by reduced physical models. *Physical Review Accelerators and Beams*, 26(8):084601. https://journals.aps.org/prab/abstract/10.1103/PhysRevAccelBeams.26.084601.

Ferri J, Siminos E, Fülöp T (2019a) Enhanced target normal sheath acceleration using colliding laser pulses. *Communications Physics*, 2(1):40. https://www.nature.com/articles/s42005-019-0140-x.

Ferri J, Siminos E, Fülöp T (2019b) Enhanced target normal sheath acceleration using colliding laser pulses. *Communications Physics*, 2(1). http://dx.doi.org/10.1038/s42005-019-0140-x.

Fiúza F, Stockem A, Boella E, Fonseca R, Silva L, Haberberger D, Tochitsky S, Gong C, Mori WB, Joshi C (2012) Laser-driven shock acceleration of monoenergetic ion beams. *Physical Review Letters*, 109(21):215001. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.109.215001.

Flacco A, Bayart E, Giaccaglia C, Monzac J, Romagnani L, Cavallone M, Patriarca A, DeMarzi L, Fouillade C, Heinrich S, Lamarre-Jouenne I, Parodi K, Rösch T, Schreiber J, Tischendorf L (2024) *Laser-FLASH : radiobiology at high dose, ultra-high dose-rate, single pulse laser-driven proton source*. https://arxiv.org/abs/2410.05086.

Flippo KA, Workman J, Gautier DC, Letzring S, Johnson RP, Shimada T (2008) Scaling laws for energetic ions from the commissioning of the new Los Alamos National Laboratory 200 TW Trident laser. *Review of Scientific Instruments*, 79(10). https://pubs.aip.org/aip/rsi/article/79/10/10E534/1070882/Scaling-laws-for-energetic-ions-from-the.

Forslund D, Freidberg J (1971) Theory of laminar collisionless shocks. *Physical Review Letters*, 27(18):1189. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.27.1189.

Forslund D, Kindel J, Lee K (1977) Theory of hot-electron spectra at high laser intensity. *Physical Review Letters*, 39(5):284. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.39.284.

Fourmaux S, Buffechoux S, Albertazzi B, Capelli D, Lévy A, Gnedyuk S, Lecherbourg L, Lassonde P, Payeur S, Antici P et al. (2013) Investigation of laser-driven proton acceleration using ultrashort, ultra-intense laser pulses. *Physics of Plasmas*, 20(1). https://pubs.aip.org/aip/pop/article/20/1/013110/1016520.

Fourmaux S, Corde S, Phuoc KT, Lassonde P, Lebrun G, Payeur S, Martin F, Sebban S, Malka V, Rousse A et al. (2011a) Single shot phase contrast imaging using laser-produced Betatron x-ray beams. *Optics letters*, 36(13):2426–2428. https://opg.optica.org/ol/fulltext.cfm?uri=ol-36-13-2426&id=219064.

Fourmaux S, Hallin E, Krol A, Bourgade J, Kieffer J (2020) X-ray phase contrast imaging of spherical capsules. *Optics Express*, 28(9):13978–13990. https://opg.optica.org/oe/fulltext.cfm?uri=oe-28-9-13978&id=431195.

Fourmaux S, Kieffer J, Hallin E (2018) Laser wakefield driven x-ray sources in Canada : a brillant future for agriculture and global food security. *Advances in X-Ray/EUV Optics and Components XIII*, SPIE, volume 10760, pages 68–74. https://www.spiedigitallibrary.org/conference-proceedings-of-spie/10760/107600E/Laser-wakefield-driven-x-ray-sources-in-Canada--a/10.1117/12.2319887.short.

Fourmaux S, Payeur S, Alexandrov A, Serbanescu C, Martin F, Ozaki T, Kudryashov A, Kieffer J (2008) Laser beam wavefront correction for ultra high intensities with the 200 TW laser system at the advanced laser light source. *Optics express*, 16(16):11987–11994. https://opg.optica.org/oe/fulltext.cfm?uri=oe-16-16-11987&id=170130.

Fourmaux S, Payeur S, Buffechoux S, Lassonde P, St-Pierre C, Martin F, Kieffer J (2011b) Pedestal cleaning for high laser pulse contrast ratio with a 100 TW class laser system. *Optics Express*, 19(9):8486–8497. https://opg.optica.org/oe/fulltext.cfm?uri=oe-19-9-8486&id=212817.

Frazier PI (2018) A tutorial on Bayesian optimization. *arXiv preprint arXiv :1807.02811*. https://arxiv.org/abs/1807.02811.

Froula D, Clayton C, Döppner T, Marsh K, Barty C, Divol L, Fonseca R, Glenzer S, Joshi C, Lu W et al. (2009) Measurements of the Critical Power for Self-Injection of Electrons in a Laser Wakefield Accelerator. *Physical Review Letters*, 103(21):215006. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.103.215006.

Fuchs J, Antici P, d'Humières E, Lefebvre E, Borghesi M, Brambrink E, Cecchetti C, Kaluza M, Malka V, Manclossi M et al. (2006) Laser-driven proton scaling laws and new paths towards energy increase. *Nature physics*, 2(1):48–54. https://www.nature.com/articles/nphys199.

Fujii T, Oishi Y, Nayuki T, Takizawa Y, Nemoto K, Kayoiji T, Horioka K, Okano Y, Hironaka Y, Nakamura KG et al. (2003) MeV-order proton and carbon ion acceleration by irradiation of 60 fs TW laser pulses on thin copper tape. *Applied physics letters*, 83(8):1524–1526. https://pubs.aip.org/aip/apl/article/83/8/1524/509216/MeV-order-proton-and-carbon-ion-acceleration-by.

Gahn C, Tsakiris G, Pukhov A, Meyer-ter Vehn J, Pretzler G, Thirolf P, Habs D, Witte K (1999) Multi-MeV electron beam generation by direct laser acceleration in high-density plasma channels. *Physical Review Letters*, 83(23):4772. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.83.4772.

Gao Y, Bin J, Haffa D, Kreuzer C, Hartmann J, Speicher M, Lindner FH, Ostermayr TM, Hilz P, Rösch TF et al. (2017) An automated, 0.5 Hz nano-foil target positioning system for intense laser plasma experiments. *High Power Laser Science and Engineering*, 5:e12. https://doi.org/10.1017/hpl.2017.10.

Garten M, Bulanov SS, Hakimi S, Obst-Huebl L, Mitchell CE, Schroeder CB, Esarey E, Geddes CG, Vay JL, Huebl A (2024) Laser-plasma ion beam booster based on hollow-channel magnetic vortex acceleration. *Physical Review Research*, 6(3):033148. https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.6.033148.

Gauthier M, Kim J, Curry C, Aurand B, Gamboa E, Göde S, Goyon C, Hazi A, Kerr S, Pak A et al. (2016) High-intensity laser-accelerated ion beam produced from cryogenic micro-jet target. *Review of Scientific Instruments*, 87(11). https://pubs.aip.org/aip/rsi/article/87/11/11D827/363050.

Geddes C, Nakamura K, Plateau GR, Toth C, Cormier-Michel E, Esarey E, Schroeder C, Cary J, Leemans WP (2008) Plasma-density-gradient injection of low absolute-momentum-spread electron bunches. *Physical review letters*, 100(21):215004. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.100.215004.

Geddes C, Toth C, Van Tilborg J, Esarey E, Schroeder C, Bruhwiler D, Nieter C, Cary J, Leemans W (2004) High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding. *Nature*, 431(7008):538–541. https://www.nature.com/articles/nature02900.

George K, Morrison J, Feister S, Ngirmang G, Smith J, Klim A, Snyder J, Austin D, Erbsen W, Frische K et al. (2019) High-repetition-rate (kHz) targets and optics from liquid microjets for high-intensity laser–plasma interactions. *High Power Laser Science and Engineering*, 7:e50. https://doi.org/10.1017/hpl.2019.35.

Geulig LD, Obst-Huebl L, Nakamura K, Bin J, Ji Q, Steinke S, Snijders AM, Mao JH, Blakely EA, Gonsalves AJ, Bulanov S, van Tilborg J, Schroeder CB, Geddes CGR, Esarey E, Roth M, Schenkel T (2022) Online charge measurement for petawatt laser-driven ion acceleration. *Review of Scientific Instruments*, 93(10). http://dx.doi.org/10.1063/5.0096423.

Ghigo A, Petrarca M, Cialdi S, Vicario C (2007) Laser temporal pulse shaping based on the DAZZLER. CARE.

Gibbon P (2005) *Short pulse laser interactions with matter : an introduction*. World Scientific. https://www.worldscientific.com/worldscibooks/10.1142/p116?srsltid=AfmBOoqr1yRGuvXSdjcr6iAtGOEzRid7xF4kx8ZwBKlagfnqYOW9BUhZ#t=aboutBook.

Giesen A, Speiser J (2007) Fifteen years of work on thin-disk lasers : results and scaling laws. *IEEE Journal of selected topics in quantum electronics*, 13(3):598–609. https://ieeexplore.ieee.org/abstract/document/4244423.

Giuffrida L, Schillaci F, Grepl F, Tryus M, Stancek S, Velyhan A, Istokskaia V, Petringa G, Cirrone G, Cupal J et al. (2023) Commissioning experiments at the ELIMAIA user beamline (Conference Presentation). *Laser Acceleration of Electrons, Protons, and Ions VII*, SPIE, PC125790S pages. https://doi.org/10.1117/12.2665598.

Giulietti A, Macchi A, Schifano E, Biancalana V, Danson C, Giulietti D, Gizzi L, Willi O (1999) Stimulated Brillouin backscattering from underdense expanding plasmas in a regime of strong

filamentation. *Physical Review E*, 59(1):1038. https://journals.aps.org/pre/abstract/10.1103/PhysRevE.59.1038.

Giulietti D, Breschi E, Galimberti M, Giulietti A, Gizzi L, Koester P, Labate L, Tomassini P, Martin P, Ceccotti T et al. (2007) High brightness laser induced multi-MeV electron/proton sources. *International Journal of Modern Physics A*, 22(22):3810–3825. https://www.worldscientific.com/doi/abs/10.1142/S0217751X07037445.

Gizzi LA, Boella E, Labate L, Baffigi F, Bilbao PJ, Brandi F, Cristoforetti G, Fazzi A, Fulgentini L, Giove D et al. (2021) Enhanced laser-driven proton acceleration via improved fast electron heating in a controlled pre-plasma. *Scientific reports*, 11(1):13728. https://www.nature.com/articles/s41598-021-93011-3.

Glenzer SH, Redmer R (2009) X-ray Thomson scattering in high energy density plasmas. *Reviews of Modern Physics*, 81(4):1625–1663. https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.81.1625.

Gonoskov A, Wallin E, Polovinkin A, Meyerov I (2019) Employing machine learning for theory validation and identification of experimental conditions in laser-plasma physics. *Scientific reports*, 9(1):7043. https://www.nature.com/articles/s41598-019-43465-3.

Gonsalves A, Nakamura K, Daniels J, Benedetti C, Pieronek C, De Raadt T, Steinke S, Bin J, Bulanov S, Van Tilborg J et al. (2019) Petawatt laser guiding and electron beam acceleration to 8 GeV in a laser-heated capillary discharge waveguide. *Physical review letters*, 122(8):084801. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.122.084801.

Green J, Murphy C, Booth N, Dance R, Gray R, MacLellan D, McKenna P, Rusby D, Wilson L (2014a) Single shot, temporally and spatially resolved measurements of fast electron dynamics using a chirped optical probe. *Journal of Instrumentation*, 9(03):P03003. https://iopscience.iop.org/article/10.1088/1748-0221/9/03/P03003/meta.

Green JS, Robinson APL, Booth N, Carroll DC, Dance RJ, Gray RJ, MacLellan DA, McKenna P, Murphy CD, Rusby D, Wilson L (2014b) High efficiency proton beam generation through target thickness control in femtosecond laser-plasma interactions. *Applied Physics Letters*, 104(21):214101. https://doi.org/10.1063/1.4879641.

Gremillet L, Bonnaud G, Amiranoff F (2002) Filamented transport of laser-generated relativistic electrons penetrating a solid target. *Physics of Plasmas*, 9(3):941–948. https://pubs.aip.org/aip/pop/article/9/3/941/265019/Filamented-transport-of-laser-generated.

Grieser S, Aurand B, Aktan E, Bonaventura D, Büscher M, Cerchez M, Engin I, Leßmann L, Mannweiler C, Prasad R, Willi O, Khoukaz A (2019) Nm-sized cryogenic hydrogen clusters for a laser-driven proton source. *Review of Scientific Instruments*, 90(4):043301. https://doi.org/10.1063/1.5080011.

Gwynne D, Kar S, Doria D, Ahmed H, Cerchez M, Fernandez J, Gray R, Green J, Hanton F, MacLellan D et al. (2014) Modified Thomson spectrometer design for high energy, multi-species ion sources. *Review of Scientific Instruments*, 85(3). https://pubs.aip.org/aip/rsi/article/85/3/033304/380784.

Ha D, Schmidhuber J (2018) World models. *arXiv preprint arXiv :1803.10122*. https://arxiv.org/abs/1803.10122.

Haberberger D, Tochitsky S, Fiuza F, Gong C, Fonseca RA, Silva LO, Mori WB, Joshi C (2012) Collisionless shocks in laser-produced plasma generate monoenergetic high-energy proton beams. *Nature Physics*, 8(1):95–99. https://www.nature.com/articles/nphys2130.

Hadjisolomou P, Tsygvintsev I, Sasorov P, Gasilov V, Korn G, Bulanov S (2020) Preplasma effects on laser ion generation from thin foil targets. *Physics of Plasmas*, 27(1). https://pubs.aip.org/aip/pop/article/27/1/013107/263126.

Hafizi B, Ting A, Sprangle P, Hubbard R (2000) Relativistic focusing and ponderomotive channeling of intense laser beams. *Physical Review E*, 62(3):4120. https://journals.aps.org/pre/abstract/10.1103/PhysRevE.62.4120.

Hafz NA, Jeong TM, Choi IW, Lee SK, Pae KH, Kulagin VV, Sung JH, Yu TJ, Hong KH, Hosokai T et al. (2008) Stable generation of GeV-class electron beams from self-guided laser–plasma channels. *Nature Photonics*, 2(9):571–577. https://www.nature.com/articles/nphoton.2008.155.

Haines MG, Wei MS, Beg FN, Stephens RB (2009) Hot-Electron Temperature and Laser-Light Absorption in Fast Ignition. *Phys. Rev. Lett.*, 102:045008. https://link.aps.org/doi/10.1103/PhysRevLett.102.045008.

Harres K, Schollmeier M, Brambrink E, Audebert P, Blažević A, Flippo K, Gautier D, Geißel M, Hegelich B, Nürnberg F et al. (2008) Development and calibration of a Thomson parabola with microchannel plate for the detection of laser-accelerated MeV ions. *Review of Scientific Instruments*, 79(9). https://pubs.aip.org/aip/rsi/article/79/9/093306/377809/Development-and-calibration-of-a-Thomson-parabola.

Hastie T, Tibshirani R, Friedman J (2008) Random forests. *The elements of statistical learning : Data mining, inference, and prediction*, Springer, pages 587–604. https://link.springer.com/chapter/10.1007/978-0-387-84858-7_15.

Hatfield PW, Gaffney JA, Anderson GJ, Ali S, Antonelli L, Başeğmez du Pree S, Citrin J, Fajardo M, Knapp P, Kettle B et al. (2021) The data-driven future of high-energy-density physics. *Nature*, 593(7859):351–361. https://www.nature.com/articles/s41586-021-03382-w.

Hegelich BM, Albright B, Cobble J, Flippo K, Letzring S, Paffett M, Ruhl H, Schreiber J, Schulze R, Fernández J (2006) Laser acceleration of quasi-monoenergetic MeV ion beams. *Nature*, 439(7075):441–444. https://www.nature.com/articles/nature04400.

Helton JC, Davis FJ (2003) Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems. *Reliability Engineering & System Safety*, 81(1):23–69. https://www.sciencedirect.com/science/article/pii/S0951832003000589.

Henig A, Kiefer D, Markey K, Gautier D, Flippo K, Letzring S, Johnson R, Shimada T, Yin L, Albright B et al. (2009) Enhanced laser-driven ion acceleration in the relativistic transparency regime. *Physical review letters*, 103(4):045002. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.103.045002.

Higginson A, Gray R, King M, Dance R, Williamson S, Butler N, Wilson R, Capdessus R, Armstrong C, Green J et al. (2018) Near-100 MeV protons via a laser-driven transparency-enhanced hybrid acceleration scheme. *Nature communications*, 9(1):724. https://www.nature.com/articles/s41467-018-03063-9.

Higginson DP, Vassura L, Gugiu MM, Antici P, Borghesi M, Brauckmann S, Diouf C, Green A, Palumbo L, Petrascu H, Sofia S, Stardubtsev M, Willi O, Kar S, Negoita F, Fuchs J (2015) Temporal Narrowing of Neutrons Produced by High-Intensity Short-Pulse Lasers. *Phys. Rev. Lett.*, 115:054802. https://link.aps.org/doi/10.1103/PhysRevLett.115.054802.

Hoffmeister G, Bellei C, Harres K, Ivanov D, Kraus D, Pelka A, Rethfeld B, Schaumann G, Roth M (2013) Influence of fs-laser desorption on target normal sheath accelerated ions. *Physical Review Special Topics—Accelerators and Beams*, 16(4):041304. https://journals.aps.org/prab/abstract/10.1103/PhysRevSTAB.16.041304.

Holmes RB (2022) Adaptive optics for directed energy : fundamentals and methodology. *AIAA Journal*, 60(10):5633–5644. https://arc.aiaa.org/doi/full/10.2514/1.J061766.

Hora H (1969) Self-focusing of laser beams in a plasma by ponderomotive forces. *Zeitschrift für Physik A Hadrons and nuclei*, 226(2):156–159. https://link.springer.com/article/10.1007/bf01392018.

Humbird KD, Peterson JL, Spears B, McClarren RG (2019) Transfer learning to model inertial confinement fusion experiments. *IEEE Transactions on Plasma Science*, 48(1):61–70. https://ieeexplore.ieee.org/abstract/document/8932676?casa_token=mGI4u5oV6psAAAAA:LPef_v_bAUW12ZTy_tV7gGS-WNSN2OgxZ6PmXVaBxxDRRMt0q_9qt3NSLe_o_Rn3h5vF1ZA-.

Hussein AE, Senabulya N, Ma Y, Streeter M, Kettle B, Dann SJ, Albert F, Bourgeois N, Cipiccia S, Cole JM et al. (2019) Laser-wakefield accelerators for high-resolution X-ray imaging of complex microstructures. *Scientific reports*, 9(1):3249. https://www.nature.com/articles/s41598-019-39845-4.

Hutter F, Hoos HH, Leyton-Brown K (2011) Sequential model-based optimization for general algorithm configuration. *Learning and intelligent optimization : 5th international conference, LION 5, rome, Italy, January 17-21, 2011. selected papers 5*, Springer, pages 507–523. https://link.springer.com/chapter/10.1007/978-3-642-25566-3_40.

Iaconis C, Walmsley IA (1998) Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses. *Optics letters*, 23(10):792–794. https://opg.optica.org/ol/abstract.cfm?uri=ol-23-10-792.

International Committee on Ultrahigh Intensity Lasers (2020) *2020 Ultrahigh Intensity Laser Facilities*. https://icuil.org. Image from ICUIL public communication material. Map of global laser facilities.

Irshad F, Karsch S, Döpp A (2023) Multi-objective and multi-fidelity Bayesian optimization of laser-plasma acceleration. *Physical Review Research*, 5(1):013063. https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.5.013063.

Ivanov K, Shulyapov S, Ksenofontov P, Tsymbalov I, Volkov R, Savel'Ev A, Brantov A, Bychenkov VY, Turinge A, Lapik A et al. (2014) Comparative study of amplified spontaneous emission and short pre-pulse impacts onto fast electron generation at sub-relativistic femtosecond laser-plasma interaction. *Physics of Plasmas*, 21(9). https://pubs.aip.org/aip/pop/article/21/9/093110/624840.

Jalas S, Kirchen M, Braun C, Eichner T, Gonzalez JB, Hübner L, Hülsenbusch T, Messner P, Palmer G, Schnepp M, Werle C, Winkler P, Leemans WP, Maier AR (2023) Tuning curves for a laser-plasma accelerator. *Phys. Rev. Accel. Beams*, 26:071302. https://link.aps.org/doi/10.1103/PhysRevAccelBeams.26.071302.

Jalas S, Kirchen M, Messner P, Winkler P, Hübner L, Dirkwinkel J, Schnepp M, Lehe R, Maier AR (2021) Bayesian optimization of a laser-plasma accelerator. *Physical review letters*, 126(10): 104801. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.126.104801.

Jiang H, Lv C, Feng K, Jiang K, Liu X, Luan S, Liu J, Wang W, Li R (2025) Parallel Bayesian optimization of free-electron lasers based on laser wakefield accelerators. *Plasma Physics and Controlled Fusion*, 67(2):025031. https://iopscience.iop.org/article/10.1088/1361-6587/adab1d/meta.

Jinno S, Fukuda Y, Sakaki H, Yogo A, Kanasaki M, Kondo K, Faenov AY, Skobelev IY, Pikuz TA, Boldarev AS, Gasilov VA (2013) Characterization of submicron-sized CO2 clusters formed with a supersonic expansion of a mixed-gas using a three-staged nozzle. *Applied Physics Letters*, 102(16):164103. https://doi.org/10.1063/1.4802915.

Jirka M, Klimo O, Matys M (2021) Relativistic plasma aperture for laser intensity enhancement. *Physical Review Research*, 3(3):033175. https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.3.033175.

Jones DR, Schonlau M, Welch WJ (1998) Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13(4):455–492. https://link.springer.com/article/10.1023/A:1008306431147.

Joyce DB, Schmid F (2010) Progress in the growth of large scale Ti : sapphire crystals by the heat exchanger method (HEM) for petawatt class lasers. *Journal of crystal growth*, 312(8):1138–1141. https://www.sciencedirect.com/science/article/pii/S0022024809009920?via%3Dihub.

Jullien A, Albert O, Burgy F, Hamoniaux G, Rousseau JP, Chambaret JP, Augé-Rochereau F, Chériaux G, Etchepare J, Minkovski N et al. (2005) 10^{-10} temporal contrast for femtosecond ultraintense lasers by cross-polarized wave generation. *Optics letters*, 30(8):920–922. https://opg.optica.org/viewmedia.cfm?r=1&rwjcode=ol&uri=ol-30-8-920&html=true.

Jung R, Tümmler J, Nubbemeyer T, Will I (2016) Thin-disk ring amplifier for high pulse energy. *Optics express*, 24(5):4375–4381. https://opg.optica.org/oe/fulltext.cfm?uri=oe-24-5-4375&id=336591.

Kain V, Hirlander S, Goddard B, Velotti FM, Della Porta GZ, Bruchon N, Valentino G (2020) Sample-efficient reinforcement learning for CERN accelerator control. *Physical Review Accelerators and Beams*, 23(12):124801. https://journals.aps.org/prab/abstract/10.1103/PhysRevAccelBeams.23.124801.

Kaloyan M, Ghazaryan S, Constantin C, Dorst R, Heuer P, Pilgram J, Schaeffer D, Niemann C (2021) Raster Thomson scattering in large-scale laser plasmas produced at high repetition rate. *Review of Scientific Instruments*, 92(9). https://pubs.aip.org/aip/rsi/article/92/9/093102/790016.

Kaloyan M, Ghazaryan S, Tripathi SP, Gekelman W, Valle MJ, Seo B, Niemann C (2022) First results from the Thomson scattering diagnostic on the large plasma device. *Instruments*, 6(2):17. https://www.mdpi.com/2410-390X/6/2/17.

Kaluza M, Schreiber J, Santala MI, Tsakiris GD, Eidmann K, Meyer-ter Vehn J, Witte KJ (2004) Influence of the laser prepulse on proton acceleration in thin-foil experiments. *Physical review letters*, 93(4):045003. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.93.045003.

Kanno I, Kunisawa T, Suzuki T, Kotera H (2007) Development of deformable mirror composed of piezoelectric thin films for adaptive optics. *IEEE journal of selected topics in quantum electronics*, 13(2):155–161. https://ieeexplore.ieee.org/abstract/document/4159979.

Karsch S, Osterhoff J, Popp A, Rowlands-Rees T, Major Z, Fuchs M, Marx B, Hörlein R, Schmid K, Veisz L et al. (2007) GeV-scale electron acceleration in a gas-filled capillary discharge waveguide. *New Journal of Physics*, 9(11):415. https://iopscience.iop.org/article/10.1088/1367-2630/9/11/415/meta.

Kaw P (2017) Nonlinear laser–plasma interactions. *Reviews of Modern Plasma Physics*, 1:1–42. https://doi.org/10.1007/s41614-017-0005-2.

Kazasidis O, Verpoort S, Soloviev O, Vdovin G, Verhaegen M, Wittrock U (2018) Extended-image-based correction of aberrations using a deformable mirror with hysteresis. *Optics Express*, 26(21):27161–27178. https://opg.optica.org/oe/fulltext.cfm?uri=oe-26-21-27161&id=398820.

Keppler S, Elkina N, Becker G, Hein J, Hornung M, Mäusezahl M, Rödel C, Tamer I, Zepf M, Kaluza M (2022) Intensity scaling limitations of laser-driven proton acceleration in the TNSA-regime. *Physical Review Research*, 4(1):013065. https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.4.013065.

Khan I, Saxena V (2024) *TNSA based proton acceleration by two oblique laser pulses in the presence of an axial magnetic field*. https://arxiv.org/abs/2406.08821.

Kim A, Dey I, Bespaly A, Komm P, Shaham A, Papeer J, Botton M, Zigler A (2024) Investigation of Pre-Pulse Effects on Ultrashort-Pulse Laser Interaction with Structured Targets. *Applied Sciences*, 15(1):237. https://www.mdpi.com/2076-3417/15/1/237.

Kim HT, Pae KH, Cha HJ, Kim IJ, Yu TJ, Sung JH, Lee SK, Jeong TM, Lee J (2013) Enhancement of electron energy to the multi-GeV regime by a dual-stage laser-wakefield accelerator pumped by petawatt laser pulses. *Physical review letters*, 111(16):165002. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.111.165002.

Kim HT, Pathak V, Hong Pae K, Lifschitz A, Sylla F, Shin JH, Hojbota C, Lee SK, Sung JH, Lee HW et al. (2017) Stable multi-GeV electron accelerator driven by waveform-controlled PW laser pulses. *Scientific reports*, 7(1):10203. https://www.nature.com/articles/s41598-017-09267-1.

Kim IJ, Pae KH, Choi IW, Lee CL, Kim HT, Singhal H, Sung JH, Lee SK, Lee HW, Nickles PV et al. (2016) Radiation pressure acceleration of protons to 93 MeV with circularly polarized petawatt laser pulses. *Physics of Plasmas*, 23(7). https://pubs.aip.org/aip/pop/article/23/7/070701/594614.

Kim J, Wilks S, Kemp A, Sherlock M, Ma T, Beg F, Mariscal D (2022) Efficient ion acceleration by multistaged intense short laser pulses. *Physical Review Research*, 4(3):L032003. https://doi.org/10.1103/PhysRevResearch.4.L032003.

King M, Wilson R, Bacon EF, Dolier EJ, Frazer TP, Goodman J, Gray RJ, McKenna P (2023) Perspectives on laser-plasma physics in the relativistic transparency regime. *The European Physical Journal A*, 59(6):132. https://link.springer.com/article/10.1140/epja/s10050-023-01043-2.

Kitagawa Y, Matsumoto T, Minamihata T, Sawai K, Matsuo K, Mima K, Nishihara K, Azechi H, Tanaka K, Takabe H et al. (1992) Beat-wave excitation of plasma wave and observation of accelerated electrons. *Physical review letters*, 68(1):48. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.68.48.

Klimo O, Psikal J, Limpouch J, Tikhonchuk VT (2008) Monoenergetic ion beams from ultrathin foils irradiated by ultrahigh-contrast circularly polarized laser pulses. *Physical Review Special Topics - Accelerators and Beams*, 11(3):031301. https://journals.aps.org/prab/abstract/10.1103/PhysRevSTAB.11.031301.

Kober J, Bagnell JA, Peters J (2013) Reinforcement learning in robotics : A survey. *The International Journal of Robotics Research*, 32(11):1238–1274. https://journals.sagepub.com/doi/full/10.1177/0278364913495721.

Kojima S, Inoue S, Dinh TH, Hasegawa N, Mori M, Sakaki H, Yamamoto Y, Sasaki T, Shiokawa K, Kondo K et al. (2020) Compact Thomson parabola spectrometer with variability of energy range and measurability of angular distribution for low-energy laser-driven accelerated ions. *Review of Scientific Instruments*, 91(5). https://pubs.aip.org/aip/rsi/article/91/5/053305/359927/Compact-Thomson-parabola-spectrometer-with.

Kokorowski S (1979) Analysis of adaptive optical elements made from piezolectric bimorphs. *JOSA*, 69(1):181–187. https://opg.optica.org/josa/fulltext.cfm?uri=josa-69-1-181&id=74734.

Kostyukov I, Kiselev S, Pukhov A (2003) X-ray generation in an ion channel. *Physics of Plasmas*, 10(12):4818–4828. https://pubs.aip.org/aip/pop/article/10/12/4818/1014286/X-ray-generation-in-an-ion-channel.

Krausz F, Ivanov M (2009) Attosecond physics. *Reviews of modern physics*, 81(1):163–234. https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.81.163.

Kruer WL, Estabrook K (1985) J$\times$ B heating by very intense laser light. *The Physics of fluids*, 28(1):430–432. https://pubs.aip.org/aip/pfl/article/28/1/430/819458/J-B-heating-by-very-intense-laser-light.

Kumar R, Sakawa Y, Döhl LN, Woolsey N, Morace A (2019) Enhancement of collisionless shock ion acceleration by electrostatic ion two-stream instability in the upstream plasma. *Physical Review Accelerators and Beams*, 22(4):043401. https://journals.aps.org/prab/abstract/10.1103/PhysRevAccelBeams.22.043401.

Kurz T, Heinemann T, Gilljohann M, Chang Y, Couperus Cabadağ J, Debus A, Kononenko O, Pausch R, Schöbel S, Assmann R et al. (2021) Demonstration of a compact plasma accelerator powered by laser-accelerated electron beams. *Nature communications*, 12(1):2895. https://www.nature.com/articles/s41467-021-23000-7.

Laserlab-Europe Consortium (2025) *Laserlab-Europe – The Integrated Initiative of European Laser Research Infrastructures*. https://www.laserlab-europe.eu/.

LaserNetUS (2025) *LaserNetUS Facilities*. https://lasernetus.org/facilities.

LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *nature*, 521(7553):436–444. https://www.nature.com/articles/nature14539.

Lécz Z, Singh P, Ter-Avetisyan S (2022) Threshold target thickness in high-contrast laser-driven ion acceleration. *Physics of Plasmas*, 29(10). https://pubs.aip.org/aip/pop/article/29/10/103104/2848055.

Ledingham K, McKenna P, McCanny T, Shimizu S, Yang J, Robson L, Zweit J, Gillies JM, Bailey J, Chimon G et al. (2004) High power laser production of short-lived isotopes for positron emission tomography. *Journal of Physics D : Applied Physics*, 37(16):2341. https://iopscience.iop.org/article/10.1088/0022-3727/37/16/019/meta.

Lee JH, Lee YC, Kang EC (2006) A cooled deformable bimorph mirror for a high power laser. *Journal of the Optical Society of Korea*, 10(2):57–62. https://opg.optica.org/josk/abstract.cfm?uri=josk-10-2-57.

Leemans W, Gonsalves A, Mao HS, Nakamura K, Benedetti C, Schroeder C, Tóth C, Daniels J, Mittelberger D, Bulanov S et al. (2014) Multi-GeV electron beams from capillary-discharge-guided subpetawatt laser pulses in the self-trapping regime. *Physical review letters*, 113(24):245002. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.113.245002.

Leemans WP, Nagler B, Gonsalves AJ, Tóth C, Nakamura K, Geddes CG, Esarey E, Schroeder CB, Hooker SM (2006) GeV electron beams from a centimetre-scale accelerator. *Nature physics*, 2(10):696–699. https://www.nature.com/articles/nphys418.

Lefaudeux N, Levecq X, Escolano L, Theis S (2013) Large deformable mirrors for beam control of high brightness lasers. 8600. https://www.spiedigitallibrary.org/conference-proceedings-of-spie/8600/1/Large-deformable-mirrors-for-beam-control-of-high-brightness-lasers/10.1117/12.2001841.short.

Lei X, Wang S, Yan H, Liu W, Dong L, Yang P, Xu B (2012) Double-deformable-mirror adaptive optics system for laser beam cleanup using blind optimization. *Optics express*, 20(20):22143–22157. https://opg.optica.org/oe/fulltext.cfm?uri=oe-20-20-22143&id=241430.

Lelievre R, Catrix E, Vallières S, Fourmaux S, Allaoua A, Anthonippillai V, Antici P, Ducasse Q, Fuchs J (2024) High repetition-rate 0.5 Hz broadband neutron source driven by the Advanced Laser Light Source. *Physics of Plasmas*, 31(9). https://pubs.aip.org/aip/pop/article/31/9/093106/3313959.

Lepetit L, Chériaux G, Joffre M (1995) Linear techniques of phase measurement by femtosecond spectral interferometry for applications in spectroscopy. *Journal of the Optical Society of America B*, 12(12):2467–2474. https://opg.optica.org/josab/abstract.cfm?uri=josab-12-12-2467.

Levy A, Ceccotti T, Popescu H, Reau F, D'Oliveira P, Monot P, Martin P, Geindre JP, Lefebvre E (2008) Proton Acceleration With High-Intensity Laser Pulses in Ultrahigh Contrast Regime. *Plasma Science, IEEE Transactions on*, 36:1808 – 1811. https://ieeexplore.ieee.org/abstract/document/4598911.

Lian CW, Ji Y, Yan R, Li J, Wang LF, Ding YK, Zheng J (2025) Two-plasmon-decay instability stimulated by dual laser beams in inertial confinement fusion. *Matter and Radiation at Extremes*, 10(1). https://pubs.aip.org/aip/mre/article/10/1/017403/3328182.

Liaw A, Wiener M et al. (2002) Classification and regression by randomForest. *R news*, 2(3):18–22. https://api.semanticscholar.org/CorpusID:3093707.

Lin J, Ma Y, Schwartz R, Woodbury D, Nees J, Mathis M, Thomas A, Krushelnick K, Milchberg H (2019) Adaptive control of laser-wakefield accelerators driven by mid-IR laser pulses. *Optics Express*, 27(8):10912–10923. https://opg.optica.org/oe/fulltext.cfm?uri=oe-27-8-10912&id=408244.

Lin J, Qian Q, Murphy J, Hsu A, Hero A, Ma Y, Thomas AG, Krushelnick K (2021) Beyond optimization—supervised learning applications in relativistic laser-plasma experiments. *Physics of Plasmas*, 28(8). https://pubs.aip.org/aip/pop/article/28/8/083102/107099.

Liu YL, Chen YC, Jao CS, Wong MS, Huang CH, Chen HW, Isayama S, Kuramitsu Y (2024) Deep learning approaches for modeling laser-driven proton beams via phase-stable acceleration. *Physics of Plasmas*, 31(1). https://pubs.aip.org/aip/pop/article/31/1/013106/3105961.

Lorenz S, Grittani G, Chacon-Golcher E, Lazzarini C, Limpouch J, Nawaz F, Nevrkla M, Vilanova L, Levato T (2019) Characterization of supersonic and subsonic gas targets for laser wakefield electron acceleration experiments. *Matter and Radiation at Extremes*, 4(1). https://pubs.aip.org/aip/mre/article/4/1/015401/252906.

Loughran B, Streeter MJ, Ahmed H, Astbury S, Balcazar M, Borghesi M, Bourgeois N, Curry C, Dann S, DiIorio S et al. (2023) Automated control and optimization of laser-driven ion acceleration. *High Power Laser Science and Engineering*, 11:e35. https://www.cambridge.org/core/journals/high-power-laser-science-and-engineering/article/automated-control-and-optimization-of-laserdriven-ion-acceleration/067E7D12CC7461E51E51B426BC7BDC40.

Lu H, Liu M, Wang W, Wang C, Liu J, Deng A, Xu J, Xia C, Li W, Zhang H et al. (2011) Laser wakefield acceleration of electron beams beyond 1 GeV from an ablative capillary discharge waveguide. *Applied Physics Letters*, 99(9). https://pubs.aip.org/aip/apl/article/99/9/091502/987000.

Lundh O, Lindau F, Persson A, Wahlström CG, McKenna P, Batani D (2007) Influence of shock waves on laser-driven proton acceleration. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 76(2):026404. https://journals.aps.org/pre/abstract/10.1103/PhysRevE.76.026404.

Ma B, Ren J, Liu L, Wei W, Chen B, Zhang S, Xu H, Hu Z, Li F, Wang X, Yin S, Feng J, Zhou X, Gao Y, Li Y, Shi X, Li J, Ren X, Xu Z, Deng Z, Qi W, Wang S, Fan Q, Cui B, Wang W, Yuan Z, Teng J, Wu Y, Cao Z, Zhao Z, Gu Y, Cao L, Zhu S, Cheng R, Lei Y, Wang Z, Zhou Z, Xiao G, Zhao H, Hoffmann DHH, Zhou W, Zhao Y (2023) *Charge equilibration of Laser-accelerated Carbon Ions in Foam Target*. https://arxiv.org/abs/2310.01047.

Macchi A, Benedetti C (2010) Ion acceleration by radiation pressure in thin and thick targets. *Nuclear Instruments and Methods in Physics Research Section A : Accelerators, Spectrometers, Detectors and Associated Equipment*, 620(1):41–45. https://www.sciencedirect.com/science/article/pii/S0168900210001257?casa_token=aOfTl3TCg_4AAAAA:YlKRrZS7LfsfFpUMUtIVBKWm6xGKYmBWWYT_xXGr90XDB0_Nj418lCvQmQv47AqfVFxuVCiJtw.

Macchi A, Borghesi M, Passoni M (2013) Ion acceleration by superintense laser-plasma interaction. *Reviews of Modern Physics*, 85(2):751–793. https://doi.org/10.1103/RevModPhys.85.751.

Macchi A (2017) A review of laser-plasma ion acceleration. *arXiv preprint arXiv :1712.06443*. https://arxiv.org/abs/1712.06443.

Mackinnon A, Patel P, Town R, Edwards M, Phillips T, Lerner S, Price D, Hicks D, Key M, Hatchett S et al. (2004) Proton radiography as an electromagnetic field and density perturbation diagnostic. *Review of Scientific Instruments*, 75(10):3531–3536. https://pubs.aip.org/aip/rsi/article/75/10/3531/350240/Proton-radiography-as-an-electromagnetic-field-and.

Mackinnon A, Sentoku Y, Patel P, Price D, Hatchett S, Key M, Andersen C, Snavely R, Freeman R (2002) Enhancement of proton acceleration by hot-electron recirculation in thin foils irradiated by ultraintense laser pulses. *Physical review letters*, 88(21):215006. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.88.215006.

Măgureanu A, Dincă L, Jalbă C, Andrei R, Burducea I, Ghiţă D, Nastasa V, Gugiu M, Asavei T, Budrigă O et al. (2022) Target characteristics used in laser-plasma acceleration of protons based on the TNSA mechanism. *Frontiers in Physics*, 10:727718. https://www.frontiersin.org/journals/physics/articles/10.3389/fphy.2022.727718/full.

Mahajan VN (1983) Strehl ratio for primary aberrations in terms of their aberration variance. *Journal of the Optical Society of America*, 73(6):860–861.

Maksimchuk A, Gu S, Flippo K, Umstadter D, Bychenkov VY (2000) Forward ion acceleration in thin films driven by a high-intensity laser. *Physical Review Letters*, 84(18):4108. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.84.4108.

Malka V et al. (2013) Review of laser wakefield accelerators. *Proceedings of IPAC2013*, pages 11–15. https://scholar.google.fr/scholar?hl=fr&as_sdt=0%2C5&q=Review+of+Laser+WakeField+Accelerators+Malka&btnG=.

Malka V, Faure J, Gauduel YA, Lefebvre E, Rousse A, Phuoc KT (2008) Principles and applications of compact laser-plasma accelerators. *Nature physics*, 4(6):447–453. https://www.nature.com/articles/nphys966.

Malka V, Faure J, Marques J, Amiranoff F, Rousseau JP, Ranc S, Chambaret J, Najmudin Z, Walton B, Mora P et al. (2001) Characterization of electron beams produced by ultrashort (30 fs) laser pulses. *Physics of Plasmas*, 8(6):2605–2608. https://pubs.aip.org/aip/pop/article/8/6/2605/730622/Characterization-of-electron-beams-produced-by.

Malka V, Fritzler S, Lefebvre E, d'Humières E, Ferrand R, Grillon G, Albaret C, Meyroneinc S, Chambaret JP, Antonetti A et al. (2004) Practicability of protontherapy using compact laser systems. *Medical physics*, 31(6):1587–1592. https://aapm.onlinelibrary.wiley.com/doi/abs/10.1118/1.1747751.

Mangles SP, Murphy C, Najmudin Z, Thomas AGR, Collier J, Dangor AE, Divall E, Foster P, Gallacher J, Hooker C et al. (2004) Monoenergetic beams of relativistic electrons from intense laser–plasma interactions. *Nature*, 431(7008):535–538. https://www.nature.com/articles/nature02939.

Margarone D, Velyhan A, Dostal J, Ullschmied J, Perin J, Chatain D, Garcia S, Bonnay P, Pisarczyk T, Dudzak R et al. (2016) Proton acceleration driven by a nanosecond laser from a cryogenic thin solid-hydrogen ribbon. *Physical Review X*, 6(4):041030. https://journals.aps.org/prx/abstract/10.1103/PhysRevX.6.041030.

Mariscal D, Djordjevic B, Anirudh R, Jayaraman-Thiagarajan J, Grace E, Simpson R, Swanson K, Galvin T, Mittelberger D, Heebner J et al. (2024) Toward machine-learning-assisted PW-class high-repetition-rate experiments with solid targets. *Physics of Plasmas*, 31(7). https://pubs.aip.org/aip/pop/article/31/7/073105/3303419.

Marquès JR, Lancia L, Loiseau P, Forestier-Colleoni P, Tarisien M, Atukpor E, Bagnoud V, Brabetz C, Consoli F, Domange J et al. (2024) Collisionless shock acceleration of protons in a plasma slab produced in a gas jet by the collision of two laser-driven hydrodynamic shockwaves. *Matter and radiation at extremes*, 9(2). https://pubs.aip.org/aip/mre/article/9/2/024001/2930959/Collisionless-shock-acceleration-of-protons-in-a.

Mathot S, Anelli G, Atieh S, Bilton A, Bulat B, Callamand T, Calvo S, Favre G, Geisser JM, Gerardin A, Grudiev A, Lombardi A, Montesinos E, Motschmann F, Pommerenke H, Richerot P, Scibor K, Timmins M, Vretenar M, Taccetti F, Benetti F, Castelli L, Chiari M, Czelusniak C, Falciano S, Fedi M, Mandò P, Manetti M, Matacotta C, Previtali E, Ruberto C, Virgili V, Giuntini L (2019) The CERN PIXE-RFQ, a transportable proton accelerator for the machina project. *Nuclear Instruments and Methods in Physics Research Section B : Beam Interactions with Materials and Atoms*, 459:153–157. https://www.sciencedirect.com/science/article/pii/S0168583X19305865.

McClarren RG, Ryu D, Drake RP, Grosskopf M, Bingham D, Chou CC, Fryxell B, Van der Holst B, Holloway JP, Kuranz CC et al. (2011) A physics informed emulator for laser-driven radiating shock simulations. *Reliability Engineering & System Safety*, 96(9):1194–1207. https://www.sciencedirect.com/science/article/pii/S0951832011000718.

McKenna P, Lindau F, Lundh O, Neely D, Persson A, Wahlström CG (2006) High-intensity laser-driven proton acceleration : influence of pulse contrast. *Philosophical Transactions of the Royal Society A : Mathematical, Physical and Engineering Sciences*, 364(1840):711–723. https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2005.1733.

McQueen CJ, Wilson R, Frazer TP, King M, Alderton M, Bacon EF, Dolier EJ, Dzelzainis T, Patel JK, P. Peat M et al. (2025) A neural network-based synthetic diagnostic of laser-accelerated proton energy spectra. *Communications Physics*, 8(1):66. https://www.nature.com/articles/s42005-025-01984-8.

Metternich M, Nazary H, Schumacher D, Brabetz C, Kroll F, Brack FE, Ehret M, Blažević A, Schramm U, Bagnoud V et al. (2022) Analyzing the filamentation of MeV-range proton bunches in a laser-driven ion beamline and optimizing their peak intensity. *Physical Review Accelerators and Beams*, 25(11):111301. https://journals.aps.org/prab/abstract/10.1103/PhysRevAccelBeams.25.111301.

Mirani F, Maffini A, Casamichiela F, Pazzaglia A, Formenti A, Dellasega D, Russo V, Vavassori D, Bortot D, Huault M, Zeraouli G, Ospina V, Malko S, Apiñaniz JI, Pérez-Hernández JA, Luis DD, Gatti G, Volpe L, Pola A, Passoni M (2021) Integrated quantitative PIXE analysis and EDX spectroscopy using a laser-driven particle source. *Science Advances*, 7(3):eabc8660. https://www.science.org/doi/abs/10.1126/sciadv.abc8660.

Mitchell TM (1997) *Machine learning*. volume 1. McGraw-hill New York. https://www.cs.cmu.edu/~tom/mlbook.html.

Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G et al. (2015) Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533. https://www.nature.com/articles/nature14236.

Mockus J (2005) The Bayesian approach to global optimization. *System Modeling and Optimization : Proceedings of the 10th IFIP Conference New York City, USA, August 31–September 4, 1981*, Springer, pages 473–481. https://doi.org/10.1007/978-94-009-0909-0.

Modena A, Najmudin Z, Dangor A, Clayton C, Marsh K, Joshi C, Malka V, Darrow C, Danson C, Neely D et al. (1995) Electron acceleration from the breaking of relativistic plasma waves. *nature*, 377(6550):606–608. https://www.nature.com/articles/377606a0.

Moore C, Ting A, Krushelnick K, Esarey E, Hubbard R, Hafizi B, Burris H, Manka C, Sprangle P (1997) Electron trapping in self-modulated laser wakefields by Raman backscatter. *Physical Review Letters*, 79(20):3909. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.79.3909.

Mora P (2003) Plasma expansion into a vacuum. *Physical Review Letters*, 90(18):185002. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.90.185002.

Mora P (2005) Thin-foil expansion into a vacuum. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 72(5):056401. https://journals.aps.org/pre/abstract/10.1103/PhysRevE.72.056401.

Moreau J, d'Humières E, Nuter R, Tikhonchuk V (2017) Stimulated Raman scattering in the relativistic regime in near-critical plasmas. *Physical Review E*, 95(1):013208. https://journals.aps.org/pre/abstract/10.1103/PhysRevE.95.013208.

Morrison J, Storm M, Chowdhury E, Akli K, Feldman S, Willis C, Daskalova R, Growden T, Berger P, Ditmire T et al. (2012) Selective deuteron production using target normal sheath acceleration. *Physics of Plasmas*, 19(3). https://pubs.aip.org/aip/pop/article/19/3/030707/918489.

Morrison J, Willis C, Freeman R, Van Woerkom L (2011) Design of and data reduction from compact Thomson parabola spectrometers. *Review of Scientific Instruments*, 82(3). https://pubs.aip.org/aip/rsi/article/82/3/033506/359441.

Morrison JT, Feister S, Frische KD, Austin DR, Ngirmang GK, Murphy NR, Orban C, Chowdhury EA, Roquemore W (2018) MeV proton acceleration at kHz repetition rate from ultra-intense laser liquid interaction. *New Journal of Physics*, 20(2):022001. https://iopscience.iop.org/article/10.1088/1367-2630/aaa8d1/meta.

Morrissey F, Fan T, Miller D, Rand D (2017) Picosecond kilohertz-class cryogenically cooled multistage Yb-doped chirped pulse amplifier. *Optics Letters*, 42(4):707–710. https://opg.optica.org/ol/fulltext.cfm?uri=ol-42-4-707&id=359473.

Mourou G, Tajima T (2012) Exploring fundamental physics at the highest-intensitylaser frontier. *URL http ://spie. org/x88664. xml*. https://spie.org/news/4221-exploring-fundamental-physics-at-the-highest-intensity-laser-frontier.

Murphy KP (2012) *Machine learning : a probabilistic perspective*. MIT press. https://books.google.fr/books?hl=fr&lr=&id=RC43AgAAQBAJ&oi=fnd&pg=PR7&dq=Murphy,+K.+P.+(2012).+Machine+Learning:+A+Probabilistic+Perspective.+MIT+Press.&ots=unjAfxKvZ8&sig=puDuZVePt33SJhllNi5vvgI63PE&redir_esc=y#v=onepage&q=Murphy%2C%20K.%20P.%20(2012).%20Machine%20Learning%3A%20A%20Probabilistic%20Perspective.%20MIT%20Press.&f=false.

Nakajima K, Fisher D, Kawakubo T, Nakanishi H, Ogata A, Kato Y, Kitagawa Y, Kodama R, Mima K, Shiraga H et al. (1995) Observation of ultrahigh gradient electron acceleration by a self-modulated intense short laser pulse. *Physical review letters*, 74(22):4428. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.74.4428.

Nakamura T, Bulanov SV, Esirkepov TZ, Kando M (2010) High-energy ions from near-critical density plasmas via magnetic vortex acceleration. *Physical review letters*, 105(13):135002. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.105.135002.

Naseri N, Bochkarev S, Rozmus W (2010) Self-channelling of relativistic laser pulses in large-scale underdense plasmas. *Physics of Plasmas*, 17(3). https://pubs.aip.org/aip/pop/article/17/3/033107/106294.

Nayuki T, Oishi Y, Fujii T, Nemoto K, Kayoiji T, Okano Y, Hironaka Y, Nakamura KG, Kondo Ki, Ueda Ki (2003) Thin tape target driver for laser ion accelerator. *Review of scientific instruments*, 74(7):3293–3296. https://pubs.aip.org/aip/rsi/article/74/7/3293/348008/Thin-tape-target-driver-for-laser-ion-accelerator.

Nedbailo R, Park J, Hollinger R, Wang S, Mariscal D, Morrison J, Song H, Zeraouli G, Scott G, Ma T et al. (2023) Compact high repetition rate Thomson parabola ion spectrometer. *Review of Scientific Instruments*, 94(2). https://doi.org/10.1063/5.0101859.

Neely D, Foster P, Robinson A, Lindau F, Lundh O, Persson A, Wahlström CG, McKenna P (2006) Enhanced proton beams from ultrathin targets driven by high contrast laser pulses. *Applied Physics Letters*, 89(2). https://pubs.aip.org/aip/apl/article/89/2/021502/902481/Enhanced-proton-beams-from-ultrathin-targets.

Nemoto K, Maksimchuk A, Banerjee S, Flippo K, Mourou G, Umstadter D, Bychenkov VY (2001) Laser-triggered ion acceleration and table top isotope production. *Applied Physics Letters*, 78(5):595–597. https://doi.org/10.1063/1.1343845.

Nishiuchi M, Daito I, Ikegami M, Daido H, Mori M, Orimo S, Ogura K, Sagisaka A, Yogo A, Pirozhkov A et al. (2009) Focusing and spectral enhancement of a repetition-rated, laser-driven, divergent multi-MeV proton beam using permanent quadrupole magnets. *Applied Physics Letters*, 94(6). https://doi.org/10.1063/1.3078291.

Nishiuchi M, Dover N, Hata M, Sakaki H, Kondo K, Lowe H, Miyahara T, Kiriyama H, Koga J, Iwata N et al. (2020) Dynamics of laser-driven heavy-ion acceleration clarified by ion charge states. *Physical Review Research*, 2(3):033081. https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.2.033081.

Niu K, Tian C (2022) Zernike polynomials and their applications. *Journal of Optics*, 24(12):123001. https://iopscience.iop.org/article/10.1088/2040-8986/ac9e08/meta.

Noaman ul Haq M, Ahmed H, Sokollik T, Fang Y, Ge X, Yuan X, Chen L (2018) Parametric scalings of laser driven protons using a high repetition rate tape drive target system. *Nuclear Instruments*

and Methods in Physics Research Section A : Accelerators, Spectrometers, Detectors and Associated Equipment, 909:164–167. https://www.sciencedirect.com/science/article/pii/S0168900218302961. 3rd European Advanced Accelerator Concepts workshop (EAAC2017).

Noaman ul Haq M, Ahmed H, Sokollik T, Yu L, Liu Z, Yuan X, Yuan F, Mirzaie M, Ge X, Chen L et al. (2017) Statistical analysis of laser driven protons using a high-repetition-rate tape drive target system. *Physical Review Accelerators and Beams*, 20(4):041301. https://journals.aps.org/prab/abstract/10.1103/PhysRevAccelBeams.20.041301.

Noll RJ (1976) Zernike polynomials and atmospheric turbulence. *Journal of the Optical Society of America*, 66(3):207–211. https://opg.optica.org/josa/abstract.cfm?uri=josa-66-3-207.

Nürnberg F, Schollmeier M, Brambrink E, Blažević A, Carroll D, Flippo K, Gautier D, Geißel M, Harres K, Hegelich B et al. (2009) Radiochromic film imaging spectroscopy of laser-accelerated proton beams. *Review of scientific instruments*, 80(3). https://pubs.aip.org/aip/rsi/article-abstract/80/3/033301/351715/Radiochromic-film-imaging-spectroscopy-of-laser?redirectedFrom=fulltext.

Obst-Huebl L, Göde S, Rehwald M, Brack FE, Branco J, Bock S, Bussmann M, Cowan T, Curry C, Fiuza F, Gauthier M, Gebhardt R, Helbig U, Huebl A, Hübner U, Irman A, Kazak L, Kim J, Kluge T, Zeil K (2017) Efficient laser-driven proton acceleration from cylindrical and planar cryogenic hydrogen jets. *Scientific Reports*, 7. https://www.nature.com/articles/s41598-017-10589-3.

Obst-Huebl L, Ziegler T, Brack FE, Branco J, Bussmann M, Cowan TE, Curry CB, Fiuza F, Garten M, Gauthier M et al. (2018) All-optical structuring of laser-driven proton beam profiles. *Nature communications*, 9(1):5292. https://www.nature.com/articles/s41467-018-07756-z.

Ohland J, Lebas N, Deo V, Guyon O, Mathieu F, Audebert P, Papadopoulos D (2025) Apollon Real-Time Adaptive Optics : astronomy-inspired wavefront stabilization in ultraintense lasers. *High Power Laser Science and Engineering*, 13:e29. https://doi.org/10.1017/hpl.2025.16.

Oishi Y, Nayuki T, Fujii T, Takizawa Y, Wang X, Yamazaki T, Nemoto K, Kayoiji T, Sekiya T, Horioka K et al. (2005) Dependence on laser intensity and pulse duration in proton acceleration by irradiation of ultrashort laser pulses on a Cu foil target. *Physics of plasmas*, 12(7). https://pubs.aip.org/aip/pop/article/12/7/073102/262458.

Osman F, Castillo R, Hora H (1999) Relativistic and ponderomotive self-focusing at laser–plasma interaction. *Journal of plasma physics*, 61(2):263–273. https://www.cambridge.org/core/journals/journal-of-plasma-physics/article/abs/relativistic-and-ponderomotive-selffocusing-at-laserplasma-interaction/BAFA97D0A2AC40B54E0940020E6E0363.

Osvay K, Csatári M, Ross I, Persson A, Wahlström CG (2005) On the temporal contrast of high intensity femtosecond laser pulses. *Laser and Particle Beams*, 23(3):327–332. https://www.cambridge.org/core/journals/laser-and-particle-beams/article/on-the-temporal-contrast-of-high-intensity-femtosecond-laser-pulses/C3910DC8034E0BDBD9B3D2306C93A6AD.

Oumbarek Espinos D, Rondepierre A, Zhidkov A, Pathak N, Jin Z, Huang K, Nakanii N, Daito I, Kando M, Hosokai T (2023) Notable improvements on LWFA through precise laser wavefront

tuning. *Scientific Reports*, 13(1):18466. https://www.nature.com/articles/s41598-023-45737-5.

Pak A, Kerr S, Lemos N, Link A, Patel P, Albert F, Divol L, Pollock B, Haberberger D, Froula D et al. (2018) Collisionless shock acceleration of narrow energy spread ion beams from mixed species plasmas using 1 μ m lasers. *Physical Review Accelerators and Beams*, 21(10):103401. https://journals.aps.org/prab/abstract/10.1103/PhysRevAccelBeams.21.103401.

Palaniyappan S, Hegelich BM, Wu HC, Jung D, Gautier DC, Yin L, Albright BJ, Johnson RP, Shimada T, Letzring S et al. (2012) Dynamics of relativistic transparency and optical shuttering in expanding overdense plasmas. *Nature Physics*, 8(10):763–769. https://www.nature.com/articles/nphys2390.

Pan S, Chen G, Tian Y, Qiu C, Luo S, Li L (2025) A comparative study of deep reinforcement learning algorithms in ultrafast mode-locked fiber lasers. *International Conference on Optical and Photonic Engineering (icOPEN 2024)*, SPIE, volume 13509, pages 488–493. https://www.spiedigitallibrary.org/conference-proceedings-of-spie/13509/1350926/A-comparative-study-of-deep-reinforcement-learning-algorithms-in-ultrafast/10.1117/12.3057863.short.

Papadopoulos D, Zou J, Le Blanc C, Chériaux G, Georges P, Druon F, Mennerat G, Ramirez P, Martin L, Fréneaux A et al. (2016) The Apollon 10 PW laser : experimental and theoretical investigation of the temporal characteristics. *High Power Laser Science and Engineering*, 4:e34. https://doi.org/10.1017/hpl.2016.34.

Park J, Bulanov SS, Bin J, Ji Q, Steinke S, Vay JL, Geddes CG, Schroeder CB, Leemans WP, Schenkel T et al. (2019) Ion acceleration in laser generated megatesla magnetic vortex. *Physics of Plasmas*, 26(10). https://pubs.aip.org/aip/pop/article/26/10/103108/264056.

Passoni M, Arioli F, Cialfi L, Dellasega D, Fedeli L, Formenti A, Giovannelli AC, Maffini A, Mirani F, Pazzaglia A, Tentori A, Vavassori D, Zavelani-Rossi M, Russo V (2019a) Advanced laser-driven ion sources and their applications in materials and nuclear science. *Plasma Physics and Controlled Fusion*, 62. https://iopscience.iop.org/article/10.1088/1361-6587/ab56c9/meta.

Passoni M, Bertagna L, Zani A (2010) Target normal sheath acceleration : theory, comparison with experiments and future perspectives. *New Journal of Physics*, 12(4):045012. https://iopscience.iop.org/article/10.1088/1367-2630/12/4/045012/meta.

Passoni M, Fedeli L, Mirani F (2019b) Superintense Laser-driven Ion Beam Analysis. *Scientific Reports*, 9:9202. https://www.nature.com/articles/s41598-019-45425-3.

Passoni M, Lontano M (2008) Theory of light-ion acceleration driven by a strong charge separation. *Physical review letters*, 101(11):115001. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.101.115001.

Passoni M, Zani A, Sgattoni A, Dellasega D, Macchi A, Prencipe I, Floquet V, Martin P, Liseykina T, Ceccotti T (2014) Energetic ions at moderate laser intensities using foam-based multi-layered targets. *Plasma Physics and Controlled Fusion*, 56(4):045001. https://iopscience.iop.org/article/10.1088/0741-3335/56/4/045001/meta.

Patel J, Armstrong C, Wilson R, Alderton M, Dolier E, Frazer T, Horne A, Lofrese A, Peat M, Woodward M et al. (2024) A scintillating fiber imaging spectrometer for active characterization of laser-driven proton beams. *High Power Laser Science and Engineering*, 12:e70. https://doi.org/10.1017/hpl.2024.62.

Patel P, Mackinnon A, Key M, Cowan T, Foord M, Allen M, Price D, Ruhl H, Springer fP, Stephens R (2003) Isochoric heating of solid-density matter with an ultrafast proton beam. *Physical review letters*, 91(12):125004. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.91.125004.

Peng J, Jury EC, Dönnes P, Ciurtin C (2021) Machine learning techniques for personalised medicine approaches in immune-mediated chronic inflammatory diseases : applications and challenges. *Frontiers in pharmacology*, 12:720694. https://www.frontiersin.org/journals/pharmacology/articles/10.3389/fphar.2021.720694/full.

Perry MD, Mourou G (1994) Terawatt to petawatt subpicosecond lasers. *Science*, 264(5161):917–924. https://www.science.org/doi/abs/10.1126/science.264.5161.917.

Phasics (2025) *QWLSI Wavefront Sensing Technology*. https://www.phasics.com/en/company/unique-wavefront-sensing-technology/. Accessed : 2025-06-04.

Piccione A, Berkery J, Sabbagh S, Andreopoulos Y (2020) Physics-guided machine learning approaches to predict the ideal stability properties of fusion plasmas. *Nuclear Fusion*, 60(4):046033. https://iopscience.iop.org/article/10.1088/1741-4326/ab7597/meta.

Polz J, Robinson AP, Kalinin A, Becker GA, Fraga RC, Hellwing M, Hornung M, Keppler S, Kessler A, Klöpfel D et al. (2019) Efficient laser-driven proton acceleration from a cryogenic solid hydrogen target. *Scientific reports*, 9(1):16534. https://www.nature.com/articles/s41598-019-52919-7.

Pompili R, Anania M, Bisesto F, Botton M, Chiadroni E, Cianchi A, Curcio A, Ferrario M, Galletti M, Henis Z et al. (2018) Ultrafast evolution of electric fields from high-intensity laser-matter interactions. *Scientific Reports*, 8(1):3243. https://www.nature.com/articles/s41598-018-21711-4.

Puyuelo-Valdes P, de Luis D, Hernandez J, Apiñaniz J, Curcio A, Henares J, Huault M, Pérez-Hernández J, Roso L, Gatti G et al. (2022) Implementation of a thin, flat water target capable of high-repetition-rate MeV-range proton acceleration in a high-power laser at the CLPU. *Plasma Physics and Controlled Fusion*, 64(5):054003. https://doi.org/10.1088/1361-6587/ac5643.

Puyuelo-Valdes P, Henares J, Hannachi F, Ceccotti T, Domange J, Ehret M, d'Humieres E, Lancia L, Marques J, Santos J et al. (2019a) Laser driven ion acceleration in high-density gas jets. *Laser Acceleration of Electrons, Protons, and Ions V*, SPIE, volume 11037, pages 14–21. https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11037/110370B/Laser-driven-ion-acceleration-in-high-density-gas-jets/10.1117/12.2520799.short.

Puyuelo-Valdes P, Henares J, Hannachi F, Ceccotti T, Domange J, Ehret M, d'Humieres E, Lancia L, Marquès JR, Ribeyre X et al. (2019b) Proton acceleration by collisionless shocks using a supersonic H2 gas-jet target and high-power infrared laser pulses. *Physics of Plasmas*, 26(12). https://pubs.aip.org/aip/pop/article-abstract/26/12/123109/253538/Proton-acceleration-by-collisionless-shocks-using?redirectedFrom=fulltext.

Puyuelo-Valdes P, Vallières S, Salvadori M, Fourmaux S, Payeur S, Kieffer JC, Hannachi F, Antici P (2021) Combined laser-based x-ray fluorescence and particle-induced x-ray emission for versatile multi-element analysis. *Scientific Reports*, 11(1):9998. https://www.nature.com/articles/s41598-021-86657-6.

Qiao B, Zepf M, Borghesi M, Geissler M (2009) Stable GeV ion-beam acceleration from thin foils by circularly polarized laser pulses. *Physical review letters*, 102(14):145002. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.102.145002.

Quinonero-Candela J, Rasmussen CE (2005) A unifying view of sparse approximate Gaussian process regression. *Journal of machine learning research*, 6(Dec):1939–1959. https://jmlr.org/papers/v6/quinonero-candela05a.html.

Radier C, Chalus O, Charbonneau M, Thambirajah S, Deschamps G, David S, Barbe J, Etter E, Matras G, Ricaud S et al. (2022) 10 PW peak power femtosecond laser pulses at ELI-NP. *High Power Laser Science and Engineering*, 10:e21. https://www.cambridge.org/core/journals/high-power-laser-science-and-engineering/article/10-pw-peak-power-femtosecond-laser-pulses-at-elinp/7F68E8E791BE58DAD69566A9DD53A96B.

Rasmussen CE (2003) Gaussian processes in machine learning. *Summer school on machine learning*, Springer, pages 63–71. https://link.springer.com/chapter/10.1007/978-3-540-28650-9_4.

Rawlinson JA, Islam MK, Galbraith DM (2002) Dose to radiation therapists from activation at high-energy accelerators used for conventional and intensity-modulated radiation therapy. *Medical Physics*, 29(4):598–608. https://aapm.onlinelibrary.wiley.com/doi/abs/10.1118/1.1463063.

Rea C, Granetz R, Montes K, Tinguely RA, Eidietis N, Hanson JM, Sammuli B (2018) Disruption prediction investigations using machine learning tools on DIII-D and Alcator C-Mod. *Plasma Physics and Controlled Fusion*, 60(8):084004. https://iopscience.iop.org/article/10.1088/1361-6587/aac7fe.

Reinhardt S, Granja C, Krejci F, Assmann W (2011) Test of pixel detectors for laser-driven accelerated particle beams. *Journal of Instrumentation*, 6(12):C12030. https://iopscience.iop.org/article/10.1088/1748-0221/6/12/C12030/meta.

Richardson DJ, Nilsson J, Clarkson WA (2010) High power fiber lasers : current status and future perspectives. *JOSA B*, 27(11):B63–B92. https://opg.optica.org/josab/fulltext.cfm?uri=josab-27-11-B63&id=206813.

Riedlinger J, Schwabe LT, Ibraimi Q, Pretzler G (2024) Electromagnetic Thomson parabola spectrometer for detection of fs laser-driven keV ions. *AIP Advances*, 14(8). https://pubs.aip.org/aip/adv/article/14/8/085304/3306432/Electromagnetic-Thomson-parabola-spectrometer-for.

Ritchie B (1994) Relativistic self-focusing and channel formation in laser-plasma interactions. *Physical Review E*, 50(2):R687. https://journals.aps.org/pre/abstract/10.1103/PhysRevE.50.R687.

Robinson A, Gibbon P, Zepf M, Kar S, Evans R, Bellei C (2009) Relativistically correct hole-boring and ion acceleration by circularly polarized laser pulses. *Plasma Physics and Controlled Fusion*, 51(2):024004. https://doi.org/10.1088/0741-3335/51/2/024004.

Robinson A, Trines R, Dover N, Najmudin Z (2012) Hole-boring radiation pressure acceleration as a basis for producing high-energy proton bunches. *Plasma Physics and Controlled Fusion*, 54(11):115001. https://doi.org/10.1088/0741-3335/54/11/115001.

Robinson APL, Zepf M, Kar S, Evans RG, Bellei C (2008) Radiation pressure acceleration of thin foils with circularly polarized laser pulses. *New Journal of Physics*, 10(1):013021. https://iopscience.iop.org/article/10.1088/1367-2630/10/1/013021/meta.

Robson L, Simpson P, Clarke RJ, Ledingham KW, Lindau F, Lundh O, McCanny T, Mora P, Neely D, Wahlström CG et al. (2007) Scaling of proton acceleration driven by petawatt-laser–plasma interactions. *Nature physics*, 3(1):58–62. https://www.nature.com/articles/nphys476.

Roddier F (1988) Curvature sensing and compensation : a new concept in adaptive optics. *Applied Optics*, 27(7):1223–1225. https://opg.optica.org/ao/fulltext.cfm?uri=ao-27-7-1223&id=165969.

Romagnani L, Bulanov S, Borghesi M, Audebert P, Gauthier J, Löwenbrück K, Mackinnon A, Patel P, Pretzler G, Toncian T et al. (2008) Observation of collisionless shocks in laser-plasma experiments. *Physical review letters*, 101(2):025004. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.101.025004.

Romagnani L, Fuchs J, Borghesi M, Antici P, Audebert P, Ceccherini F, Cowan T, Grismayer T, Kar S, Macchi A et al. (2005) Dynamics of electric fields driving the laser acceleration of multi-MeV protons. *Physical review letters*, 95(19):195001. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.95.195001.

Ross IN, Matousek P, New GH, Osvay K (2002) Analysis and optimization of optical parametric chirped pulse amplification. *Journal of the Optical Society of America B*, 19(12):2945–2956. https://doi.org/10.1364/JOSAB.19.002945.

Roth M, Brambrink E, Audebert P, Basko M, Blazevic A, Clarke R, Cobble J, Cowan T, Fernandez J, Fuchs J et al. (2005) Laser accelerated ions in ICF research prospects and experiments. *Plasma physics and controlled fusion*, 47(12B):B841. https://iopscience.iop.org/article/10.1088/0741-3335/47/12B/S66/meta.

Roth M, Cowan T, Key M, Hatchett S, Brown C, Fountain W, Johnson J, Pennington D, Snavely R, Wilks S et al. (2001) Fast ignition by intense laser-accelerated proton beams. *Physical review letters*, 86(3):436. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.86.436.

Roth M, Jung D, Falk K, Guler N, Deppert O, Devlin M, Favalli A, Fernandez J, Gautier D, Geissel M, Haight R, Hamilton CE, Hegelich BM, Johnson RP, Merrill F, Schaumann G, Schoenberg K, Schollmeier M, Shimada T, Taddeucci T, Tybo JL, Wagner F, Wender SA, Wilde CH, Wurden GA (2013) Bright Laser-Driven Neutron Source Based on the Relativistic Transparency of Solids. *Phys. Rev. Lett.*, 110:044802. https://link.aps.org/doi/10.1103/PhysRevLett.110.044802.

Roth M, Schollmeier M (2016) Ion acceleration-target normal sheath acceleration. *Proceedings of the 2014 CAS-CERN Accelerator School : Plasma Wake Acceleration*. https://e-publishing.cern.ch/index.php/CYR/article/view/222.

Rousse A, Phuoc KT, Shah R, Pukhov A, Lefebvre E, Malka V, Kiselev S, Burgy F, Rousseau JP, Umstadter D et al. (2004) Production of a keV X-Ray Beam from Synchrotron Radiation in

Relativistic Laser-Plasma Interaction. *Physical review letters*, 93(13):135005. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.93.135005.

Roussel R, Edelen AL, Boltz T, Kennedy D, Zhang Z, Ji F, Huang X, Ratner D, Garcia AS, Xu C et al. (2024) Bayesian optimization algorithms for accelerator physics. *Physical review accelerators and beams*, 27(8):084801. https://journals.aps.org/prab/abstract/10.1103/PhysRevAccelBeams.27.084801.

Ruskov R, Bingham R, Silva L, Harper M, Aboushelbaya R, Myatt J, Norreys P (2024) Statistical theory of the broadband two-plasmon decay instability. *Journal of Plasma Physics*, 90(6). http://dx.doi.org/10.1017/S0022377824000953.

Sadowski M, Czaus K, Malinowski K, Skladnik-Sadowska E, Zebrowski J (2009) Mass-and energy-analyses of ions from plasma by means of a miniature Thomson spectrometer. *Review of Scientific Instruments*, 80(5). https://pubs.aip.org/aip/rsi/article/80/5/053504/282370.

Sakabe S, Mochizuki T, Yamanaka T, Yamanaka C (1980) Modified Thomson parabola ion spectrometer of wide dynamic range. *Review of Scientific Instruments*, 51(10):1314–1315. https://pubs.aip.org/aip/rsi/article/51/10/1314/309809/Modified-Thomson-parabola-ion-spectrometer-of-wide.

Salcedo A, Focia R, Ram A, Bers A (2003) Studies of stimulated Raman backscattering and associated nonlinear laser–plasma interactions. *Nuclear fusion*, 43(12):1759. https://iopscience.iop.org/article/10.1088/0029-5515/43/12/020/meta?casa_token=PzCJhDJAZjAAAAAA:TebDYJdWsx2JFDUTcOrulTHOcjVrZ0C2hAr5k_a7CV9cpjZXq5JfiubzbzoobfWwgX0lT3jjwH6oZwgK4AJt1oQhiQ.

Samarkin V, Sheldakova J, Toporovskiy V, Rukosuev A, Kudryashov A (2021) High-spatial resolution stacked-actuator deformable mirror for correction of atmospheric wavefront aberrations. *Applied Optics*, 60(23):6719–6724. https://opg.optica.org/ao/fulltext.cfm?uri=ao-60-23-6719&id=453850.

Samarkin VV, Aleksandrov AG, Jitsuno T, Romanov P, Rukosuev AL, Kudryashov AV (2015) Study of a wide-aperture combined deformable mirror for high-power pulsed phosphate glass lasers. *Quantum Electronics*, 45(12):1086. https://iopscience.iop.org/article/10.1070/QE2015v045n12ABEH015961.

Saraceno CJ, Heckl OH, Baer CR, Südmeyer T, Keller U (2011) Pulse compression of a high-power thin disk laser using rod-type fiber amplifiers. *Optics express*, 19(2):1395–1407. https://opg.optica.org/oe/fulltext.cfm?uri=oe-19-2-1395&id=209441.

Saraceno CJ, Sutter D, Metzger T, Ahmed MA (2019) The amazing progress of high-power ultrafast thin-disk lasers. *Journal of the European Optical Society-Rapid Publications*, 15(1):15. https://jeos.edpsciences.org/articles/jeos/abs/2019/01/41476_2019_Article_108/41476_2019_Article_108.html.

Sarri G, Macchi A, Cecchetti C, Kar S, Liseykina T, Yang X, Dieckmann ME, Fuchs J, Galimberti M, Gizzi L et al. (2012) Dynamics of Self-Generated, Large Amplitude Magnetic Fields Following High-Intensity Laser Matter Interaction. *Physical review letters*, 109(20):205002. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.109.205002.

Schillaci F, Maggiore M, Velyhan A, Cirrone G, Cuttone G, Margarone D, Palumbo GP, Pisciotta P, Rifuggiato D, Romano F et al. (2014) Calibration and energy resolution study of a high dispersive power Thomson Parabola Spectrometer with monochromatic proton beams. *Journal of Instrumentation*, 9(10):T10003. https://iopscience.iop.org/article/10.1088/1748-0221/9/10/T10003.

Schlegel T, Naumova N, Tikhonchuk V, Labaune C, Sokolov I, Mourou G (2009) Relativistic laser piston model : Ponderomotive ion acceleration in dense plasmas using ultraintense laser pulses. *Physics of Plasmas*, 16(8). https://pubs.aip.org/aip/pop/article-abstract/16/8/083103/316403/Relativistic-laser-piston-model-Ponderomotive-ion?redirectedFrom=fulltext.

Schmitz B, Kreuter D, Boine-Frankenheim O, Margarone D (2023) Modeling of a liquid leaf target tnsa experiment using particle-in-cell simulations and deep learning. *Laser and Particle Beams*, 2023:e3. https://doi.org/10.1155/2023/2868112.

Schnürer M, Ter-Avetisyan S, Busch S, Risse E, Kalachnikov M, Sandner W, Nickles P (2005) Ion acceleration with ultrafast laser driven water droplets. *Laser and Particle Beams*, 23(3):337–343. https://www.cambridge.org/core/journals/laser-and-particle-beams/article/ion-acceleration-with-ultrafast-laser-driven-water-droplets/0E1B633619B05F582EBA8651EAB45FD7.

Schreiber J, Bell F, Grüner F, Schramm U, Geissler M, Schnürer M, Ter-Avetisyan S, Hegelich BM, Cobble J, Brambrink E et al. (2006) Analytical model for ion acceleration by high-intensity laser pulses. *Physical review letters*, 97(4):045005. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.97.045005.

Schulman J, Wolski F, Dhariwal P, Radford A, Klimov O (2017) Proximal policy optimization algorithms. *arXiv preprint arXiv :1707.06347*. https://arxiv.org/abs/1707.06347.

Schweighofer N, Doya K (2003) Meta-learning in reinforcement learning. *Neural Networks*, 16(1):5–9. https://www.sciencedirect.com/science/article/pii/S0893608002002289?casa_token=LozehQnURDgAAAAA:U5-XgrcRMns4y0maQIE6ii_4WlJYMDLsHTLpfS9V_uO-HcWXK5xAfn6EH-h_fEmPjcw8i6CE.

Sciscio M, Barberio M, Liberatore C, Veltri S, Laramée A, Palumbo L, Legaré F, Antici P (2017) Analysis of induced stress on materials exposed to laser-plasma radiation during high-intensity laser experiments. *Applied Surface Science*, 421:200–204. https://www.sciencedirect.com/science/article/pii/S0169433216327040.

Scott G, Carroll D, Astbury S, Clarke R, Hernandez-Gomez C, King M, Alejo A, Arteaga I, Dance R, Higginson A et al. (2018) Dual ion species plasma expansion from isotopically layered cryogenic targets. *Physical review letters*, 120(20):204801. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.120.204801.

Scuderi V, Milluzzo G, Doria D, Alejo A, Amico A, Booth N, Cuttone G, Green J, Kar S, Korn G et al. (2020) TOF diagnosis of laser accelerated, high-energy protons. *Nuclear Instruments and Methods in Physics Research Section A : Accelerators, Spectrometers, Detectors and Associated Equipment*, 978:164364. https://www.sciencedirect.com/science/article/pii/S0168900220307610.

Senthilkumaran V, Beier N, Fourmaux S, Shabaninezhad P, Stinehart J, Zhou L, Moore J, Hussein A (2024) Laser-driven betatron x rays for high-throughput imaging of additively manufactured materials. *Review of Scientific Instruments*, 95(12). https://pubs.aip.org/aip/rsi/article/95/12/123510/3324953.

Sentoku Y, Kemp AJ (2008) Numerical methods for particle simulations at extreme densities and temperatures : Weighted particles, relativistic collisions and reduced currents. *Journal of computational Physics*, 227(14):6846–6861. https://www.sciencedirect.com/science/article/pii/S0021999108001988.

Sgattoni A, Londrillo P, Macchi A, Passoni M (2012) Laser ion acceleration using a solid target coupled with a low-density layer. *Physical Review E*, 85(3). http://dx.doi.org/10.1103/PhysRevE.85.036405.

Shahriari B, Swersky K, Wang Z, Adams RP, De Freitas N (2015) Taking the human out of the loop : A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175. https://ieeexplore.ieee.org/abstract/document/7352306.

Shalloo R, Dann S, Gruse JN, Underwood C, Antoine A, Arran C, Backhouse M, Baird C, Balcazar M, Bourgeois N et al. (2020) Automation and control of laser wakefield accelerators using Bayesian optimization. *Nature communications*, 11(1):6355. https://www.nature.com/articles/s41467-020-20245-6.

Shields MD, Zhang J (2016) The generalization of Latin hypercube sampling. *Reliability Engineering & System Safety*, 148:96–108. https://www.sciencedirect.com/science/article/pii/S0951832015003543.

Siders C, Haefner C (2016) High-power lasers for science and society. Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States).

Silva LO, Marti M, Davies JR, Fonseca RA, Ren C, Tsung FS, Mori WB (2004) Proton shock acceleration in laser-plasma interactions. *Physical Review Letters*, 92(1):015002. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.92.015002.

Silver D, Huang A, Maddison CJ, Guez A, Sifre L, Van Den Driessche G, Schrittwieser J, Antonoglou I, Panneershelvam V, Lanctot M et al. (2016) Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587):484–489. https://www.nature.com/articles/nature16961.

Silver D, Hubert T, Schrittwieser J, Antonoglou I, Lai M, Guez A, Lanctot M, Sifre L, Kumaran D, Graepel T et al. (2018) A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144. https://www.science.org/doi/10.1126/science.aar6404.

Simpson R, Mariscal D, Kim J, Lemos N, Grace E, Swanson K, Scott G, Djordjevic B, Ma T (2023) Investigation of boosted proton energies through proton radiography of target normal sheath acceleration fields in the multi-ps regime. *Physics of Plasmas*, 30(10). https://pubs.aip.org/aip/pop/article-abstract/30/10/103103/2914211/Investigation-of-boosted-proton-energies-through?redirectedFrom=fulltext.

Simpson R, Mariscal D, Kim J, Scott G, Williams G, Grace E, McGuffey C, Wilks S, Kemp A, Lemos N et al. (2021a) Demonstration of TNSA proton radiography on the National Ignition Facility

Advanced Radiographic Capability (NIF-ARC) laser. *Plasma Physics and Controlled Fusion*, 63(12):124006. https://doi.org/10.1088/1361-6587/ac2349.

Simpson RA, Scott G, Mariscal D, Rusby D, King P, Grace E, Aghedo A, Pagano I, Sinclair M, Armstrong C et al. (2021b) Scaling of laser-driven electron and proton acceleration as a function of laser pulse duration, energy, and intensity in the multi-picosecond regime. *Physics of Plasmas*, 28(1). https://pubs.aip.org/aip/pop/article/28/1/013108/1059993.

Smola AJ, Schölkopf B (2004) A tutorial on support vector regression. *Statistics and computing*, 14:199–222. https://link.springer.com/article/10.1023/B:STCO.0000035301.49549.88.

Snavely R, Key M, Hatchett S, Cowan T, Roth M, Phillips T, Stoyer M, Henry E, Sangster T, Singh M et al. (2000) Intense high-energy proton beams from petawatt-laser irradiation of solids. *Physical review letters*, 85(14):2945. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.85.2945.

Snoek J, Larochelle H, Adams RP (2012) Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems*, 25. https://proceedings.neurips.cc/paper/2012/hash/05311655a15b75fab86956663e1819cd-Abstract.html.

Snyder J, Morrison J, Feister S, Frische K, George K, Le M, Orban C, Ngirmang G, Chowdhury E, Roquemore W (2020) Background pressure effects on MeV protons accelerated via relativistically intense laser-plasma interactions. *Scientific Reports*, 10(1):18245. https://www.nature.com/articles/s41598-020-75061-1.

Song IU, Choi JH, Rhee HG, Kihm H, Yang HS (2024) Extended wavefront reconstruction for quadri-wave lateral shearing interferometry. *Optics and Lasers in Engineering*, 178:108212. https://www.sciencedirect.com/science/article/pii/S014381662400191X.

Speiser JL, Miller ME, Tooze J, Ip E (2019) A comparison of random forest variable selection methods for classification prediction modeling. *Expert systems with applications*, 134:93–101. https://www.sciencedirect.com/science/article/pii/S0957417419303574.

Spiceworks (2024) *What is Reinforcement Learning?* https://www.spiceworks.com/tech/artificial-intelligence/articles/what-is-reinforcement-learning/. Accessed : 2025-05-08.

Spot Intelligence (2024) *Support Vector Regression (SVR) : A Gentle Introduction*. https://spotintelligence.com/2024/05/08/support-vector-regression-svr/. Accessed : 2025-05-07.

Sprangle P, Esarey E, Ting A (1990) Nonlinear theory of intense laser-plasma interactions. *Physical review letters*, 64(17):2011. https://doi.org/10.1103/PhysRevLett.64.2011.

Sprangle P, Tang CM, Esarey E (1987) Relativistic self-focusing of short-pulse radiation beams in plasmas. *IEEE transactions on plasma science*, 15(2):145–153. https://ieeexplore.ieee.org/abstract/document/4316677?casa_token=6Cw7pascUAkAAAAA:0Qiw3tJ8cl5Uwoy135j-N_iBgsueRohVfUK8aKX3WipZO-chxRdLWM-kPvtyXAUjSGWkJBPG.

Sprangle P, Tang CM, Esarey E (2007) Relativistic self-focusing of short-pulse radiation beams in plasmas. *IEEE transactions on plasma science*, 15(2):145–153. https://ieeexplore.ieee.org/document/4316677.

Strehlow J, Kim J, Bailly-Grandvaux M, Bolaños S, Smith H, Haid A, Alfonso EL, Aniculaesei C, Chen H, Ditmire T et al. (2022) A laser parameter study on enhancing proton generation from microtube foil targets. *Scientific Reports*, 12(1):10827. https://www.nature.com/articles/s41598-022-14881-9.

Strickland D, Mourou G (1985) Compression of amplified chirped optical pulses. *Optics communications*, 55(6):447–449. https://www.sciencedirect.com/science/article/abs/pii/0030401885901518.

Su L, Liu B, Lin X, Liu F, Du F, Liu X, Zheng Y, Ge X, Li Y, Sheng Z et al. (2013) The influence of target material and thickness on proton energy and angular distribution. *Science China Physics, Mechanics and Astronomy*, 56:457–461. https://link.springer.com/article/10.1007/s11433-012-4961-9.

Südmeyer T, Kränkel C, Baer CR, Heckl O, Saraceno CJ, Golling M, Peters R, Petermann K, Huber G, Keller U (2009) High-power ultrafast thin disk laser oscillators and their potential for sub-100-femtosecond pulse generation. *Applied Physics B*, 97:281–295. https://link.springer.com/article/10.1007/s00340-009-3700-z.

Sun C, Kaiser E, Brunton SL, Kutz JN (2020) Deep reinforcement learning for optical systems : A case study of mode-locked lasers. *Machine Learning : Science and Technology*, 1(4):045013. https://iopscience.iop.org/article/10.1088/2632-2153/abb6d6/meta.

Sutton RS, Barto AG et al. (1998) *Reinforcement learning : An introduction*. volume 1. MIT press Cambridge. chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/https://web.stanford.edu/class/psych209/Readings/SuttonBartoIPRLBook2ndEd.pdf.

Taira T, Tulloch WM, Byer RL (1997) Modeling of quasi-three-level lasers and operation of cw Yb : YAG lasers. *Applied optics*, 36(9):1867–1874. https://opg.optica.org/ao/fulltext.cfm?uri=ao-36-9-1867&id=42516.

Tajima T, Dawson JM (1979) Laser electron accelerator. *Physical review letters*, 43(4):267. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.43.267.

Tajima T (2010) Laser acceleration and its future. *Proceedings of the Japan Academy, Series B*, 86(3):147–157. https://www.jstage.jst.go.jp/article/pjab/86/3/86_3_147/_article/-char/ja/.

Takagi Y, Iwata N, d'Humieres E, Sentoku Y (2021) Multivariate scaling of maximum proton energy in intense laser driven ion acceleration. *Physical Review Research*, 3(4):043140. https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.3.043140.

Tang H, Tangtartharakul K, Babjak R, Yeh I, Albert F, Chen H, Campbell P, Ma Y, Nilson P, Russell B et al. (2024) The influence of laser focusing conditions on the direct laser acceleration of electrons. *New Journal of Physics*, 26(5):053010. https://iopscience.iop.org/article/10.1088/1367-2630/ad3be4/meta.

Tata S, Mondal A, Sarkar S, Lad AD, Krishnamurthy M (2017) A gated Thomson parabola spectrometer for improved ion and neutral atom measurements in intense laser produced plasmas. *Review of Scientific Instruments*, 88(8). https://pubs.aip.org/aip/rsi/article/88/8/083305/955796.

Tazes I, Passalidis S, Kaselouris E, Mancelli D, Karvounis C, Skoulakis A, Fitilis I, Bakarezos M, Papadogiannis N, Dimitriou V et al. (2024) Efficient Magnetic Vortex Acceleration by femtosecond laser interaction with long living optically shaped gas targets in the near critical density plasma regime. *Scientific Reports*, 14(1):4945. https://www.nature.com/articles/s41598-024-54475-1.

The MathWorks Inc. (2024) *Reinforcement Learning Documentation*. https://www.mathworks.com/discovery/reinforcement-learning.html. Accessed : 2025-04-03.

Thiele I, Ferri J, Siminos E, Gremillet L, Smetanina E, Dmitriev A, Cantono G, Wahlström CG, Fülöp T (2019) *Enhancement of laser-driven ion acceleration in non-periodic nanostructured targets*. https://arxiv.org/abs/1905.11131.

Tidman DA (1971) Shock waves in collisionless plasmas. https://api.semanticscholar.org/CorpusID:123085169.

Ting A, Krushelnick K, Moore C, Burris H, Esarey E, Krall J, Sprangle P (1996) Temporal evolution of self-modulated laser wakefields measured by coherent Thomson scattering. *Physical review letters*, 77(27):5377. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.77.5377.

Toncian T, Borghesi M, Fuchs J, d'Humières E, Antici P, Audebert P, Brambrink E, Cecchetti CA, Pipahl A, Romagnani L et al. (2006) Ultrafast laser-driven microlens to focus and energy-select mega-electron volt protons. *Science*, 312(5772):410–413. https://www.science.org/doi/full/10.1126/science.1124412.

Tong SF, Sheng ZM, Yu M (2018) Target normal sheath acceleration of protons using triple-layer target with long preplasma. *Physical Review Accelerators and Beams*, 21(5):051303. https://journals.aps.org/prab/abstract/10.1103/PhysRevAccelBeams.21.051303.

Toporovskiy, Skvortsov A, Kudryashov A, Samarkin V, Sheldakova YV, Pshonkin D (2019a) Flexible bimorphic mirror with high density of control electrodes for correcting wavefront aberrations. *Journal of Optical Technology*, 86(1):32–38. https://opg.optica.org/jot/abstract.cfm?uri=jot-86-1-32.

Toporovskiy V, Kudryashov A, Samarkin V, Sheldakova J, Rukosuev A, Skvortsov A, Pshonkin D (2019b) Bimorph deformable mirror with a high density of electrodes to correct for atmospheric distortions. *Applied Optics*, 58(22):6019–6026. https://opg.optica.org/ao/fulltext.cfm?uri=ao-58-22-6019&id=416163.

Toporovskiy V, Kudryashov A, Skvortsov A, Rukosuev A, Samarkin V, Galaktionov I (2022) State-of-the-art technologies in piezoelectric deformable mirror design. *Photonics*, MDPI, volume 9, 321 pages. https://www.mdpi.com/2304-6732/9/5/321.

Torrisi L, Cutroneo M, Rosinski M, Badziak J, Parys P, Wołowski J, Zaraś-Szydłowska A, Torrisi A (2019) *Near-3-MeV protons from target-normal-sheath-acceleration femtosecond laser irradiating advanced targets*. https://onlinelibrary.wiley.com/doi/full/10.1002/ctpp.201800127.

Torrisi L, Cutroneo M, Ullschmied J (2017) TNSA and ponderomotive plasma production in enriched carbon polyethylene foils. *Physics of Plasmas*, 24(4). https://pubs.aip.org/aip/pop/article/24/4/043112/110200.

Tournois P (1997) Acousto-optic programmable dispersive filter for adaptive compensation of group delay time dispersion in laser systems. *Optics communications*, 140(4-6):245–249. https://www.sciencedirect.com/science/article/pii/S0030401897001533.

Trebino R, DeLong KW, Fittinghoff DN, Sweetser JN, Krumbügel MA, Richman BA, Kane DJ (1997) Measuring ultrashort laser pulses in the time-frequency domain using frequency-resolved optical gating. *Review of Scientific Instruments*, 68(9):3277–3295. https://pubs.aip.org/aip/rsi/article-abstract/68/9/3277/322634/Measuring-ultrashort-laser-pulses-in-the-time.

Treffert F, Curry CB, Chou HG, Crissman CJ, DePonte DP, Fiuza F, Glenn GD, Hollinger RC, Nedbailo R, Park J et al. (2022a) High-repetition-rate, multi-MeV deuteron acceleration from converging heavy water microjets at laser intensities of 1021 W/cm2. *Applied Physics Letters*, 121(7). https://doi.org/10.1063/5.0098973.

Treffert F, Glenn GD, Chou HG, Crissman C, Curry CB, DePonte DP, Fiuza F, Hartley NJ, Ofori-Okai B, Roth M et al. (2022b) Ambient-temperature liquid jet targets for high-repetition-rate HED discovery science. *Physics of Plasmas*, 29(12). https://pubs.aip.org/aip/pop/article/29/12/123105/2843606.

Trieschmann J, Vialetto L, Gergs T (2023) Machine learning for advancing low-temperature plasma modeling and simulation. *Journal of Micro/Nanopatterning, Materials, and Metrology*, 22(4): 041504–041504. https://www.spiedigitallibrary.org/journals/journal-of-micro-nanopatterning-materials-and-metrology/volume-22/issue-4/041504/Review--Machine-learning-for-advancing-low-temperature-plasma-modeling/10.1117/1.JMM.22.4.041504.full.

Tušar T, Filipič B (2014) Visualization of Pareto front approximations in evolutionary multiobjective optimization : A critical review and the prosection method. *IEEE Transactions on Evolutionary Computation*, 19(2):225–245. https://ieeexplore.ieee.org/abstract/document/6777535.

Umstadter D, Chen SY, Maksimchuk A, Mourou G, Wagner R (1996) Nonlinear optics in relativistic plasmas and laser wake field acceleration of electrons. *Science*, 273(5274):472–475. https://www.science.org/doi/abs/10.1126/science.273.5274.472.

Vallières S, Morabito A, Veltri S, Scisciò M, Barberio M, Antici P (2017) Laser-driven proton acceleration with nanostructured targets. *Laser Acceleration of Electrons, Protons, and Ions IV*, SPIE, volume 10240, pages 14–19. https://www.spiedigitallibrary.org/conference-proceedings-of-spie/10240/1024009/Laser-driven-proton-acceleration-with-nanostructured-targets/10.1117/12.2265913.short.

Vallières S, Salvadori M, Permogorov A, Cantono G, Svendsen K, Chen Z, Sun S, Consoli F, d'Humières E, Wahlström CG et al. (2021) Enhanced laser-driven proton acceleration using nanowire targets. *Scientific Reports*, 11(1):2226. https://www.nature.com/articles/s41598-020-80392-0.

Vallières S, Scisciò M, Veltri S, Barberio M, d'Humières E, Antici P (2018) Enhancement of laser-driven proton beams using nanostructured solid foils. *High Intensity Lasers and High Field Phenomena*, Optica Publishing Group, pages HM3A–6. https://opg.optica.org/abstract.cfm?uri=hilas-2018-HM3A.6.

Vallières S, Salvadori M, Puyuelo-Valdes P, Payeur S, Fourmaux S, Consoli F, Verona C, d'Humières E, Chicoine M, Roorda S, Schiettekatte F, Antici P (2020) Thomson parabola and time-of-flight detector cross-calibration methodology on the ALLS 100 TW laser-driven ion acceleration beamline. *Review of Scientific Instruments*, 91(10):103303. https://doi.org/10.1063/5.0020257.

Vargas M, Schumaker W, He Z, Behm K, Chvykov V, Hou B, Krushelnick K, Maksimchuk A, Nees J, Yanovsky V et al. (2019) X-ray phase contrast imaging of additive manufactured structures using a laser wakefield accelerator. *Plasma Physics and Controlled Fusion*, 61(5):054009. https://iopscience.iop.org/article/10.1088/1361-6587/ab0e4f/meta.

Vladisavlevici IM, Ehret M, Filippov E, García-García E, Mendez C, Ruíz MO, Varela Ó, Volpe L, Pérez-Hernández JA (2025) Theoretical Study of the Pre-Plasma Density Scale Length's Influence on the Absorption Efficiency in Laser–Solid Interaction at Relativistic Laser Intensities for PW-Class Lasers. *Photonics*, MDPI AG, volume 12, 71 pages. https://www.proquest.com/docview/3159538444?pq-origsite=gscholar&fromopenview=true&sourcetype=Scholarly%20Journals.

Vozenin MC, Bourhis J, Durante M (2022) Towards clinical translation of FLASH radiotherapy. *Nature Reviews Clinical Oncology*, 19(12):791–803. https://www.nature.com/articles/s41571-022-00697-z.

Vshivkov VA, Naumova NM, Pegoraro F, Bulanov S (1998) Nonlinear electrodynamics of the interaction of ultra-intense laser pulses with a thin foil. *Physics of Plasmas*, 5(7):2727–2741. https://pubs.aip.org/aip/pop/article-abstract/5/7/2727/444218/Nonlinear-electrodynamics-of-the-interaction-of.

Wagner F, Deppert O, Brabetz C, Fiala P, Kleinschmidt A, Poth P, Schanz V, Tebartz A, Zielbauer B, Roth M et al. (2016) Maximum proton energy above 85 mev from the relativistic interaction of laser pulses with micrometer thick ch 2 targets. *Physical review letters*, 116(20):205002. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.116.205002.

Wang W, Sun X, Sun F, Lv Z, Glize K, Shi Z, Xu Y, Zhang Z, Wu F, Hu J et al. (2025) Enhanced Proton Acceleration via Petawatt Laguerre-Gaussian Lasers. *arXiv preprint arXiv :2501.12683*. https://arxiv.org/abs/2501.12683.

Wang W (2024) Proton acceleration driven by relativistic femtosecond Laguerre–Gaussian lasers. *Reviews of Modern Plasma Physics*, 8(1):1–29. https://link.springer.com/article/10.1007/s41614-024-00174-3.

Wang X, Liu X, Lu X, Chen J, Long Y, Li W, Chen H, Chen X, Bai P, Li Y et al. (2022) 13.4 fs, 0.1 Hz OPCPA front end for the 100 PW-class laser facility. *Ultrafast Science*. https://spj.science.org/doi/10.34133/2022/9894358.

Watkins CJ, Dayan P (1992) Q-learning. *Machine learning*, 8:279–292. https://link.springer.com/article/10.1007/bf00992698.

Weber S, Riconda C, Tikhonchuk V et al. (2008) Stimulated Brillouin backscattering for laser-plasma interaction in the strong coupling regime. https://journals.aps.org/pre/abstract/10.1103/PhysRevE.59.1038.

Weichman K, Santos J, Fujioka S, Toncian T, Arefiev A (2020) Generation of focusing ion beams by magnetized electron sheath acceleration. *Scientific Reports*, 10(1):18966. https://www.nature.com/articles/s41598-020-75915-8.

Weiss K, Khoshgoftaar TM, Wang D (2016) A survey of transfer learning. *Journal of Big data*, 3:1–40. https://link.springer.com/article/10.1186/S40537-016-0043-6.

Wilks S, Kruer W, Tabak M, Langdon A (1992) Absorption of ultra-intense laser pulses. *Physical review letters*, 69(9):1383. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.69.1383.

Wilks S, Langdon A, Cowan T, Roth M, Singh M, Hatchett S, Key M, Pennington D, MacKinnon A, Snavely R (2001) Energetic proton generation in ultra-intense laser–solid interactions. *Physics of plasmas*, 8(2):542–549. https://pubs.aip.org/aip/pop/article/8/2/542/264618/Energetic-proton-generation-in-ultra-intense-laser.

Willim C, Vieira J, Malka V, Silva LO (2023) Proton acceleration with intense twisted laser light. *Physical Review Research*, 5(2):023083. https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.5.023083.

Wlodarczyk KL, Bryce E, Schwartz N, Strachan M, Hutson D, Maier RR, Atkinson D, Beard S, Baillie T, Parr-Burman P et al. (2014) Scalable stacked array piezoelectric deformable mirror for astronomy and laser processing applications. *Review of Scientific Instruments*, 85(2). https://pubs.aip.org/aip/rsi/article/85/2/024502/356462.

Würl M (2018) *On the spectrometry of laser-accelerated particle bunches and laser-driven proton radiography*. Thèse de doctorat, Ludwig-Maximilians-Universität München. https://edoc.ub.uni-muenchen.de/23381/.

Xiao Q, Pan X, Jiang Y, Wang J, Du L, Guo J, Huang D, Lu X, Cui Z, Yang S et al. (2021) High-contrast OPCPA front end in high-power petawatt laser facility based on the ps-OPCPA seed system. *Optics Express*, 29(11):15980–15994. https://doi.org/10.1364/OE.425420.

Xu F, Uszkoreit H, Du Y, Fan W, Zhao D, Zhu J (2019) Explainable AI : A brief survey on history, research areas, approaches and challenges. *Natural language processing and Chinese computing : 8th cCF international conference, NLPCC 2019, dunhuang, China, October 9–14, 2019, proceedings, part II 8*, Springer, pages 563–574. https://link.springer.com/chapter/10.1007/978-3-030-32236-6_51.

Xu L, Yu L, Liang X, Chu Y, Hu Z, Ma L, Xu Y, Wang C, Lu X, Lu H et al. (2013) High-energy noncollinear optical parametric–chirped pulse amplification in LBO at 800 nm. *Optics Letters*, 38(22):4837–4840. https://opg.optica.org/ol/fulltext.cfm?uri=ol-38-22-4837&id=274461.

Yang L, Zhang Y, Zhang Z, Li Y, Xiang Y, Dong J, Wei Y, Chang S, Lu R (2022) Effects of sample surface morphology on laser-induced breakdown spectroscopy. *Journal of Analytical Atomic Spectrometry*, 37(8):1642–1651. https://pubs.rsc.org/en/content/articlehtml/2022/ja/d2ja00101b.

Yao C, Li J, Hao L, Yan R, Wang C, Lei A, Ding Y, Zheng J (2024a) Anomalous hot electron generation from two-plasmon decay instability driven by broadband laser pulses with intensity modulations. *Nuclear Fusion*, 64(10):106013. https://iopscience.iop.org/article/10.1088/1741-4326/ad6c62/meta.

Yao W, Lelièvre R, Waltenspiel T, Cohen I, Allaoua A, Antici P, Beck A, Cohen E, Davoine X, D'humières E et al. (2024b) Enhanced energy, conversion efficiency and collimation of protons driven by high-contrast and ultrashort laser pulses. *Applied Sciences*, 14(14):6101. https://www.mdpi.com/2076-3417/14/14/6101.

Ye H, Gu Y, Zhang X, Wang S, Tan F, Zhang J, Yang Y, Yan Y, Wu Y, Huang W et al. (2022) Fast optimization for betatron radiation from laser wakefield acceleration based on Bayesian optimization. *Results in Physics*, 43:106116. https://www.sciencedirect.com/science/article/pii/S2211379722007306.

Yin L, Albright B, Hegelich B, Bowers KJ, Flippo K, Kwan T, Fernández J (2007) Monoenergetic and GeV ion acceleration from the laser breakout afterburner using ultrathin targets. *Physics of plasmas*, 14(5). https://pubs.aip.org/aip/pop/article-abstract/14/5/056706/930540/Monoenergetic-and-GeV-ion-acceleration-from-the?redirectedFrom=fulltext.

Yin L, Albright B, Hegelich B, Fernández J (2006) GeV laser ion acceleration from ultrathin targets : The laser break-out afterburner. *Laser and Particle Beams*, 24(2):291–298. https://www.cambridge.org/core/journals/laser-and-particle-beams/article/gev-laser-ion-acceleration-from-ultrathin-targets-the-laser-breakout-afterburner/AB013C8F8617E14CE09B65E69CCCA750.

Yogo A, Maeda T, Hori T, Sakaki H, Ogura K, Nishiuchi M, Sagisaka A, Kiriyama H, Okada H, Kanazawa S, Shimomura T, Nakai Y, Tanoue M, Sasao F, Bolton PR, Murakami M, Nomura T, Kawanishi S, Kondo K (2011) Measurement of relative biological effectiveness of protons in human cancer cells using a laser-driven quasimonoenergetic proton beamline. *Applied Physics Letters*, 98(5):053701. https://doi.org/10.1063/1.3551623.

Yogo A, Sato K, Nishikino M, Mori M, Teshima T, Numasaki H, Murakami M, Demizu Y, Akagi S, Nagayama S et al. (2009) Application of laser-accelerated protons to the demonstration of DNA double-strand breaks in human cancer cells. *Applied Physics Letters*, 94(18). https://pubs.aip.org/aip/apl/article/94/18/181502/336514.

Yoon JW, Jeon C, Shin J, Lee SK, Lee HW, Choi IW, Kim HT, Sung JH, Nam CH (2019) Achieving the laser intensity of 5.5× 1022 W/cm2 with a wavefront-corrected multi-PW laser. *Optics express*, 27(15):20412–20420. https://opg.optica.org/oe/viewmedia.cfm?uri=oe-27-15-20412&html=true.

Zamkotsian F, Liotard A, Lanzoni P, Conedera V, Fabre N, Camon H (2006) Electrostatic micro-deformable mirror for adaptive optics. *Advances in Adaptive Optics II*, SPIE, volume 6272, pages 676–687. https://www.spiedigitallibrary.org/conference-proceedings-of-spie/6272/627222/Electrostatic-micro-deformable-mirror-for-adaptive-optics/10.1117/12.671632.short.

Zeil K, Kraft S, Bock S, Bussmann M, Cowan T, Kluge T, Metzkes J, Richter T, Sauerbrey R, Schramm U (2010) The scaling of proton energies in ultrashort pulse laser plasma acceleration. *New Journal of Physics*, 12(4):045015. https://iopscience.iop.org/article/10.1088/1367-2630/12/4/045015/meta.

Zeng X, Zhou K, Zuo Y, Zhu Q, Su J, Wang X, Wang X, Huang X, Jiang X, Jiang D et al. (2017) Multi-petawatt laser facility fully based on optical parametric chirped-pulse amplification. *Optics letters*, 42(10):2014–2017. https://opg.optica.org/ol/fulltext.cfm?uri=ol-42-10-2014&id=366565.

Zervas MN, Codemard CA (2014) High power fiber lasers : a review. *IEEE Journal of selected topics in Quantum Electronics*, 20(5):219–241. https://ieeexplore.ieee.org/abstract/document/6808413.

Zhang T (2023) *Applying Machine Learning Methods to Laser Acceleration of Protons : Lessons Learned from Synthetic Data*. Thèse de doctorat, The Ohio State University. https://kb.osu.edu.

Zhidkov A, Nemoto K, Nayuki T, Oishi Y, Fuji T (2007) Giant electromagnetic vortex and MeV monoenergetic electrons generated by short laser pulses in underdense plasma near quarter critical density region. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 76(1): 016401. https://journals.aps.org/pre/abstract/10.1103/PhysRevE.76.016401.

Zhou O, Tsai HE, Ostermayr TM, Fan-Chiang L, Van Tilborg J, Schroeder CB, Esarey E, Geddes CG (2021) Effect of nozzle curvature on supersonic gas jets used in laser–plasma acceleration. *Physics of Plasmas*, 28(9). https://pubs.aip.org/aip/pop/article-abstract/28/9/093107/595429/Effect-of-nozzle-curvature-on-supersonic-gas-jets?redirectedFrom=fulltext.

Zhou Z, Li X, Zare RN (2017) Optimizing chemical reactions with deep reinforcement learning. *ACS central science*, 3(12):1337–1344. https://pubs.acs.org/doi/full/10.1021/acscentsci.7b00492.

Zhu Z, Li Y, Chen J, Ma J, Chu J (2017) Development of a unimorph deformable mirror with water cooling. *Optics express*, 25(24):29916–29926. https://opg.optica.org/oe/fulltext.cfm?uri=oe-25-24-29916&id=376953.

Ziegler T, Göthel I, Assenbaum S, Bernert C, Brack FE, Cowan TE, Dover NP, Gaus L, Kluge T, Kraft S et al. (2024) Laser-driven high-energy proton beams from cascaded acceleration regimes. *Nature Physics*, 20(7):1211–1216. https://www.nature.com/articles/s41567-024-02505-0.