

PAPER • OPEN ACCESS

Implementing Bayesian optimization to compare injection regimes and pulse propagation in two stage LWFA

To cite this article: Christopher Gardner and Franklin Dollar 2025 *New J. Phys.* **27** 124301

View the [article online](#) for updates and enhancements.

You may also like

- [Leveraging trust for joint multi-objective and multi-fidelity optimization](#)
Faran Irshad, Stefan Karsch and Andreas Döpp
- [Automated deep brain stimulation programming with safety constraints for tremor suppression in patients with Parkinson's disease and essential tremor](#)
Parisa Sarikhani, Benjamin Ferleger, Kyle Mitchell et al.
- [Cell viability prediction and optimization in extrusion-based bioprinting via neural network-based Bayesian optimization models](#)
Dorsa Mohammadrezaei, Lena Podina, Johanna De Silva et al.

**PAPER****OPEN ACCESS****RECEIVED**
5 August 2025**REVISED**
18 November 2025**ACCEPTED FOR PUBLICATION**
25 November 2025**PUBLISHED**
8 December 2025

Original content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Implementing Bayesian optimization to compare injection regimes and pulse propagation in two stage LWFA

Christopher Gardner*  and Franklin Dollar 

STROBE, NSF Science & Technology Center, University of California, Irvine, CA, United States of America

* Author to whom any correspondence should be addressed.

E-mail: gardnecl@uci.edu**Keywords:** laser wakefield acceleration, plasma acceleration, laser plasma interaction, particle in cell, Bayesian optimization, pulse compressionSupplementary material for this article is available [online](#)**Abstract**

Predicting the properties of electrons accelerated to relativistic energies via laser wakefield acceleration is dependent on a wide range of parameters that are difficult to efficiently optimize. Such problems would benefit from the use of Bayesian optimization schemes, which if implemented correctly may efficiently find optimums in complex parameter spaces. By separating the laser wakefield mechanism into two stages, an initially higher density plasma of high Z gases may be used to inject higher charge while a low density homogeneous plasma can be used to efficiently accelerate electrons with the same laser pulse. However, the nonlinear relationship between the two stages makes the determination of optimal beam parameters difficult. We show that through the use of a Bayesian optimization process on 2.5D3V particle-in-cell simulations we are able to efficiently map out regions of interest across a wide variety of physical regimes while keeping down the computational costs of the parameter scan. Numerically modeling the two stage system, we varied the injection stage density profile and length and examined the impact on electron beam charge, energy spectrum, and emittance. We also examined the impact on the laser pulse itself with regards to the temporal pulse duration. Our analysis identified specific regions in the parameter space where variations in pulse propagation and injection processes significantly influenced the characteristics of the resulting electron beam. This approach enabled us to efficiently identify numerous physical regimes with options on how to tune the electron beam for various applications.

1. Introduction

Laser wakefield acceleration (LWFA) is a mechanism for accelerating electrons within a plasma wave driven by an intense ultra-fast laser [1]. We define an ultra-fast laser as a laser with a pulse duration less than 100 fs. LWFA has the capability to reach acceleration gradients on the order of GV cm^{-1} [2, 3] enabling relativistic electron beams from millimeters of acceleration length. The resulting electron beams have fs time-scale durations [4] making them suitable for applications like ultra-fast electron diffraction [5]. LWFA is also widely used for sources of high brightness fs x-ray beams via the betatron mechanism [6, 7]. LWFA beams can also produce light across the x-ray and gamma ray spectrum through a variety of other mechanisms, such as Compton Scattering and Bremsstrahlung [8].

In the so-called ‘bubble regime’ of LWFA [9–13], the driving laser is so intense that the ponderomotive force can evacuate nearly all the electrons in a region behind the pulse. An ion cavity forms in the wake since the higher mass ions move on a longer time scale than the electrons. Behind this cavity an electron density wake forms as the electrons that were pushed transversely outward by the laser are pulled back by the ions. The resulting longitudinal charge separation between the wake and the ion cavity forms an accelerating structure for an electron that is trapped inside the bubble. The bubble moves

with the sub-luminal group velocity of the laser, which limits the maximum energy the electron may gain. This limit on the acceleration length is referred to as the dephasing length.

An important aspect of LWFA is how the laser parameters and propagation impacts the bubble structure. At relativistic intensities above $10^{18} \text{ W cm}^{-2}$ relativistic effects lead the bubble radius become proportional to $\sqrt{a_0}$ so that changes in the laser intensity as the pulse evolves also modify the bubble size and shape [14]. Here, $a_0 = \frac{e|\vec{E}_0|}{m_e\omega_L c}$ is the normalized vector potential, e is the elementary charge, m_e is the electron mass, $|\vec{E}_0|$ is the peak electric field of the laser, ω_L is the angular frequency of the laser, and c is the speed of light. Resonantly driving a wakefield also requires matching the laser pulse duration to the plasma wavelength. Specifically, the pulse duration (τ) must satisfy $c\tau \leq \lambda_p/2 \approx \pi c/\omega_p$ [1], where λ_p is the plasma wavelength, and $\omega_p = \sqrt{n_e e^2 / (m_e \epsilon_0)}$ is the plasma frequency in terms of the electron density (n_e), and vacuum permittivity (ϵ_0). At higher electron density, λ_p is shorter and the bubble's electron density gradients are steeper. Higher field strength bubbles therefore require shorter pulse durations. Recently, there has been success in using pulses with durations only a few optical cycles long (<10 fs at 800 nm), for LWFA [15–18]. These experiments typically use millijoule scale energy pulses externally compressed to few cycle incident on high density gas targets to produce MeV electron beams with high charge. However, scaling the external compression to few-cycle of laser drivers at Joule energies is difficult due to issues maintaining laser beam quality [19, 20]. Pulse compression can be performed *in-situ* using density gradients in the bubble itself to create a variation in the refractive index. This effect causes a pulse to self-focus [21] and self-compress [22], matching the pulse's waist and length to the bubble radius. Using this method, pulses as short as 10 fs have been measured experimentally [23].

Trapping electrons in the accelerating structure of an LWFA in a controlled way is also a non-trivial problem. Trapping may occur when an electron is both spatially inside the bubble and has a longitudinal velocity greater than the wake phase velocity. The process of getting an electron trapped in this structure is called injection. Controlling the injection of charge into the wakefield is a critical aspect of a LWFA setup since the resulting electron beam's properties such as the energy spread, charge, and emittance are sensitive to the injection process used [24–29]. Two common techniques for injecting electrons from the background plasma into the bubble are density transition injection [30], and ionization injection [31, 32]. In density transition injection, electrons are injected as the bubble expands at the transition from high to low electron density. Density transition injection is also sometimes referred to as down-ramp injection [33, 34]. During ionization injection, the laser field ionizes inner-shell electrons from a high- Z gas inside the bubble. These electrons then may be injected when they fall behind the laser into the wake [31, 32].

A higher electron density may be desirable for injecting more charge while a lower density is desirable for increasing the acceleration length to maximize energy gain. This is because at high electron density, more electrons are available and stronger fields can more readily accelerate them into the trapped region of phase space while at low density, the higher group velocity of the laser in the plasma enables a longer accelerating length. These competing interests motivate separating the gas target into two stages: one stage for injection and a subsequent stage for acceleration [35]. To enable more control over the resulting electron beam properties through staging, the injection should be localized to the injection stage. An advantage of density transition injection in a two stage setup is that it naturally localizes the injection to the location of the density transition. Similarly, by confining high- Z species to a first stage, ionization injection can also be localized [36]. This enables greater control over the resulting electron beam properties. For instance, localizing the electron injection in a higher density first stage can, under the right conditions, reduce the overall energy spread of the electron beam as more electrons take advantage of the full acceleration length in the second stage as opposed to continuous injection throughout the accelerator [28, 37]. Multi-stage accelerators have a nontrivial dependence on the initial conditions, motivating the use of numerical models for understanding and optimizing parameters [38–40]. Here, we define a 'stage' as an independent gas target that is a part of a single LWFA. This differs from schemes that use multiple laser drivers incident on a cascading chain of gas targets, each one acting as an independent LWFA [41, 42].

In this article, we present a numerical investigation and optimization of a two-stage LWFA system driven by a multi-terawatt (TW) laser where the first stage is designed to simultaneously control the electron injection and laser pulse propagation. Here, the first stage models a thin argon gas jet with a shock front created by a physical obstruction to achieve the ionization and high density requirements. The second stage is an extended homogeneous helium accelerating stage at a fixed density. The laser-plasma interaction was modeled using 2.5D3V particle-in-cell (PIC) simulations. By controlling

the density and length of the first stage, we can control the compressed laser pulse duration and self-focusing. Since the impact of the initial conditions on resulting electron and laser parameters is complicated and nonlinear, we use Gaussian process (GP) regression and Bayesian optimization on the density and length of the first stage to determine optima in a computationally efficient manner. This process enabled us to identify 4 distinct acceleration regimes in the parameter space. At lower first stage densities, we observe electron beams with over 100 pC and hundreds of MeV energy. As the density is further increased, different qualitative electron beam properties are observed depending on the thickness of the first stage, likely due to pulse compression and self-focusing. We found optimal values in each regime for increasing the accelerated charge and electron energy for varying injection processes.

2. Methods

We performed simulations on the University of California, Irvine HPC3 computing cluster, utilizing their GPU nodes. Using the PIC code WarpX [43], we systematically vary the length scale and density for the first stage which was followed by a fixed millimeter-scale helium accelerator stage. The simulations employed a quasi-cylindrical 2.5D3V geometry where the fields are decomposed into azimuthal modes with optimizations for GPUs [44]. This geometry takes advantage of the radial symmetry of both the laser field and the target density profile, significantly reducing computational resource requirements while still capturing physical phenomena that would be neglected in 1D or 2D geometries. To reduce numerical Cherenkov radiation [45] present in finite difference time domain solvers, we used a pseudo-spectral Maxwell solver [46]. The convention we will use is that the direction of propagation is the z -axis, and the radial coordinate is r with the laser polarization along the x coordinate and the y coordinate in the orthogonal direction to x and z . We define the z -axis 0 as being the location of the transition between the two stages. The laser travels from negative to positive values of z . The simulations were run with the first two azimuthal modes. The window size was $40\ \mu\text{m}$ (r) by $80\ \mu\text{m}$ (z). The mesh resolution was $\lambda/5.12$ in the radial direction and $\lambda/20.48$ in the direction of propagation. The window moved at the speed of light in the direction of propagation. Each cell has 8 neutral atom macro-particles, with ionization modeled with the ADK model [47]. The laser was a linearly polarized Gaussian beam with a $10\ \mu\text{m}$ beam waist, $0.8\ \mu\text{m}$ central wavelength (λ), 1 J energy, and a full-width at half-maximum (FWHM) pulse duration of 35 fs corresponding to a peak power of ~ 27 TW. These laser parameters are not atypical for LWFA experiments [48]. The focused field strength in free space has a normalized vector potential (a_0) of 2.84. The free space focus was fixed at $100\ \mu\text{m}$ before the transition, i.e. $z_{\text{foc}} = -100\ \mu\text{m}$.

We chose the first stage density profile to realistically achieve a high density, using an analytical expression which approximates a high density gas jet with an obstruction to form a shock front [28]. This configuration was chosen as opposed to a uniform cell due to practical limitations related to minimum length, density gradients, and maximum density. The first stage has two adjustable parameters, the maximum density and the length scale. Analytically, we represented the first stage as being defined in the region from $z_1 < z < 0$ and given as a function of the varied parameters n_0 and z_1 by the logistic function in the form shown below,

$$n_{\text{Ar}} = \left(\frac{2}{1 + e^{-K(z-z_1)}} - 1 \right) n_0. \quad (1)$$

The form of this equation makes the density start at 0 at the position z_1 , increase linearly near $z = z_1$, and eventually plateau at n_0 . The K parameter is set to make the density at the end of this stage at $0\ \mu\text{m}$ equal to $0.905 \cdot n_0$. The second stage, for acceleration, has a constant density profile meant to model an extended gas jet or cell [49]. The density and length of this stage are kept fixed. The acceleration stage extends from $0\ \mu\text{m}$ to the approximate dephasing length at 2.3 mm with a constant helium neutral density of $3 \times 10^{18}\ \text{cm}^{-3}$. A diagram of the density profile is shown in figure 1.

We used neutral argon, with density n_{Ar} , as the gas medium for the first stage. We used argon over more typical gases such as helium to improve the pulse propagation. For instance, to mitigate excessive self-focusing, argon was used rather than helium in the first stage to enable ionization-induced defocusing. This mitigates the relativistic self focusing in the plasma which enables a higher ratio of pulse compression to self focusing. Additionally, the ionization gradient in argon can also enhance the pulse compression [50]. At the laser parameters used we expect, based on a calculation of the barrier suppression ionization [51], to ionize the L and M shells of the argon on axis to get 16 electrons per initially neutral argon and to fully ionize the helium. Because the L shell electrons require higher intensities to ionize, they are ionized close to the peak intensity enabling ionization injection.

In a realistic experiment, the density profile for an injector stage would have a complicated profile depending on the details of the target design like gas intermixing between stages. Additionally, the specific density profile of the down-ramp at the transition between stages will impact the injection process [26, 27]. These details have been omitted to restrict the injection and pulse compression to a clearly defined first stage, to facilitate physical understanding.

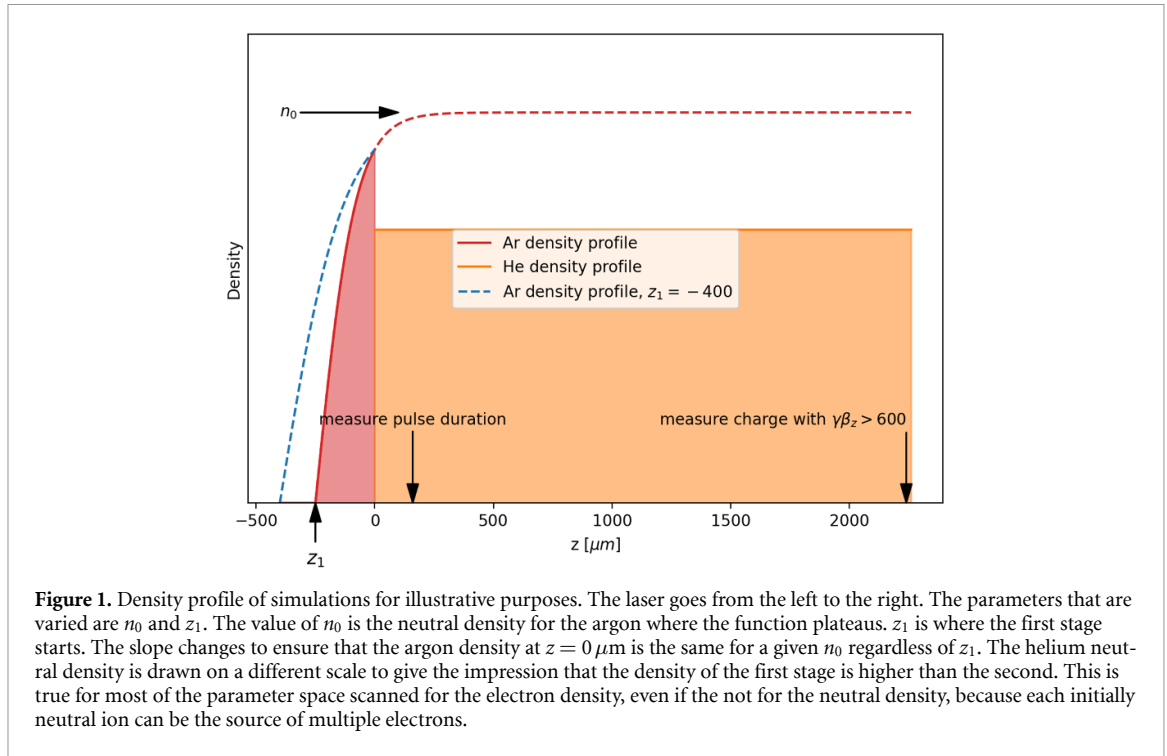
The accelerated bunch was optimized by simultaneously scanning the peak argon neutral density, n_0 , and the starting location of the first stage, z_1 and calculating the charge of the electrons with a normalized momentum satisfying $\gamma\beta_z > 600$. We choose this condition as many applications depend on having a high charge while also requiring high energy. From a preliminary simulation with only the helium stage, we found that the momentum distribution was centered on $\gamma\beta_z = 995$, so $\gamma\beta_z > 600$ was chosen to only count electrons that are accelerated to at least 60% of this value. For reference, this normalized momentum corresponds to an energy of 307 MeV. To efficiently use finite computation time, we alternated between exploring the varying parameters n_0 and z_1 randomly and optimizing the charge with $\gamma\beta_z > 600$ with Bayesian optimization using GP regression [52]. A GP regression makes it straightforward to incorporate results from anywhere in the parameter space into a statistical model based on the results. The GP regression model interpolates the results, which for a two-dimensional parameter space can then be plotted as a topographical map to visualize the model. As more simulations are run, the model becomes more accurate and can be used to predict optimal parameters with a Bayesian optimization process [53]. We used a GP model with a Matérn 5/2 kernel added to a white kernel. The Matérn kernel incorporates assumptions about the length-scale and smoothness of the model while the white kernel accounts for noise to prevent overfitting. By repeating the loop of exploring the parameter space, visualizing with GPs, and optimizing, we identified patterns and areas of interest in the results that would otherwise have been less apparent.

The parameter scanning was done through the python library Optimas [39, 54] using Bayesian optimization to find optima in the two-dimensional parameter space more efficiently and Sobol sampling to sample the entire parameter space. The GP regression shown in figure 3 was made using Optimas, while those shown in figures 2 and 5 were made using the scikit-learn package [55].

3. Results

3.1. Parameter scan and optimization

A total of 135 simulations were run to generate the data presented in this article. The median simulation took approximately 2 h to complete on a NVIDIA V100 GPU and simulations were typically run in pairs so that two GPUs could be utilized simultaneously. The first 25 simulations took initial values for parameters n_0 and z_1 from pseudo-random Sobol sampling. The next 14 simulations had Bayesian optimization enabled to find and explore optima for charge with $\gamma\beta_z > 600$. The optimization algorithm quickly found a local minimum on the bottom of the range for z_1 in the region of n_0 near $1.5 \times 10^{18} \text{ cm}^{-3}$. After this local optimum had been explored, we continued with Sobol sampling to explore more of the parameter space for simulations 41–50. In this second random sampling scan, we saw a point with a significantly higher charge of -136 pC with $\gamma\beta_z > 600$ at $n_0 = 0.53 \times 10^{18} \text{ cm}^{-3}$ and $z_1 = -319.9 \mu\text{m}$. Before this point, the charge magnitude was below 52 pC for comparison. After finding this point, Bayesian optimization was used for the next simulations from 51 to 75. However, since the hyperparameters of the underlying model for the optimization algorithm are not known *a priori*, the model treated this high charge point as an outlier. As a result, the length-scale for the kernel was too long and it continued to exploit the local minimum it had already found previously rather than explore the new optimal point. In order to encourage the optimization algorithm to explore the region around $n_0 \approx 0.53 \times 10^{18} \text{ cm}^{-3}$, the limits for the parameter space were temporarily changed to $n_0 < 0.7 \times 10^{18} \text{ cm}^{-3}$ for simulations 82–111, after which this density region was well sampled. The subsequent 4 simulations sampled from the full parameter space, and with the additional points, the algorithm was able to self-correct the earlier issue with the kernel's length-scale to find the global optimum for charge with $\gamma\beta_z > 600$ in simulation 113. This simulation is shown in the discussion section in figure 8(c). The remaining simulations were run with Sobol sampling, sometimes with manually restricted regions of the parameter space to reduce the uncertainty in regions that had been under-sampled previously. The resulting GP model for the charge with $\gamma\beta_z > 600$ is shown in figure 2(c). Although the optimization was run on the charge with $\gamma\beta_z > 600$, the results for that and other momentum cutoffs were analyzed on the same dataset and are also shown in figure 2 to display more information about the resulting momentum spectra. Electrons outside the primary bubble are excluded from the charge results in



figures 2(a)–(d) as well as in the rest of the results section by only counting charge within the rectangular position selection annotated in figure 8. This ensures that charge from secondary bubbles or that is injected late in the simulation is excluded in the analysis.

3.2. Charge results

The most charge accelerated to $\gamma\beta_z > 600$ was -142 pC when the value of n_0 was $0.52 \times 10^{18} \text{ cm}^{-3}$ and the start of the first stage was at $z_1 = -100 \mu\text{m}$. The approximate corresponding electron density at the transition between the two stages would be $n_e \approx 0.905 \cdot 16n_0 = 7.5 \times 10^{18} \text{ cm}^{-3}$, assuming that the L and M shells of argon are ionized but the K shell is not. The second helium stage would have an electron density of $6 \times 10^{18} \text{ cm}^{-3}$ giving a ratio of 1.25 between the first and second stage electron densities at the transition. This density ratio is similar to the values in the literature in studies of density transition injection such as in [26]. Around this point, it can be seen in figure 2 that there is a band of n_0 where the quantity of charge injected for $\gamma\beta_z > 400$ is maximized around $n_0 \approx 5 \times 10^{17} \text{ cm}^{-3}$. As the momentum cutoff for the reported charge is increased this band of n_0 , where the magnitude of charge is maximized becomes narrower on the n_0 axis and shifts to a lower n_0 value.

Another feature of figure 2 is that in the region nearby $n_0 \approx 2 \times 10^{18} \text{ cm}^{-3}$ and $z_1 \approx -150 \mu\text{m}$ there is a peak in the magnitude of charge for the $\gamma\beta_z > 400$ and the $\gamma\beta_z > 500$ momentum cutoff. However, for higher momentum cutoffs this feature goes away with little charge in the $\gamma\beta_z > 600$ plot and minimal charge in this region in the $\gamma\beta_z > 700$ plot.

There is a local maximum for the charge magnitude in the bottom right for the higher momentum cutoff, $\gamma\beta_z > 600$ and $\gamma\beta_z > 700$, plots in figures 2(c) and (d) around $n_0 \approx 1.6 \times 10^{18} \text{ cm}^{-3}$. The simulation in this region with the most charge with $\gamma\beta_z > 600$ was located at the parameters, $n_0 = 1.87 \times 10^{18} \text{ cm}^{-3}$ and $z_1 = -400 \mu\text{m}$ and had a charge of -50 pC above that momentum cutoff. This local maximum goes away for the lower momentum cutoffs of $\gamma\beta_z > 400$ and $\gamma\beta_z > 500$.

3.3. Maximum energy results

The contour plot in figure 3 displays a map of the GP regression of the maximum energy of the electrons. The simulation with the highest maximum energy electron had an energy of 754 MeV with $n_0 = 0.26 \times 10^{18} \text{ cm}^{-3}$ and $z_1 = -282 \mu\text{m}$. The highest maximum energies occurred for simulations on the left side of the plot with $n_0 \lesssim 0.4 \times 10^{18} \text{ cm}^{-3}$. The band of n_0 near $0.6 \times 10^{18} \text{ cm}^{-3}$ where there was a maximum for the charge magnitude in figure 2 has a decreasing maximum energy for increasing n_0 in the maximum energy plot. There is another local maximum in the maximum energy on the bottom for n_0 values between $0.5 \times 10^{18} \text{ cm}^{-3}$ and $2 \times 10^{18} \text{ cm}^{-3}$ and $z_1 \lesssim 0$. This is the same region that the

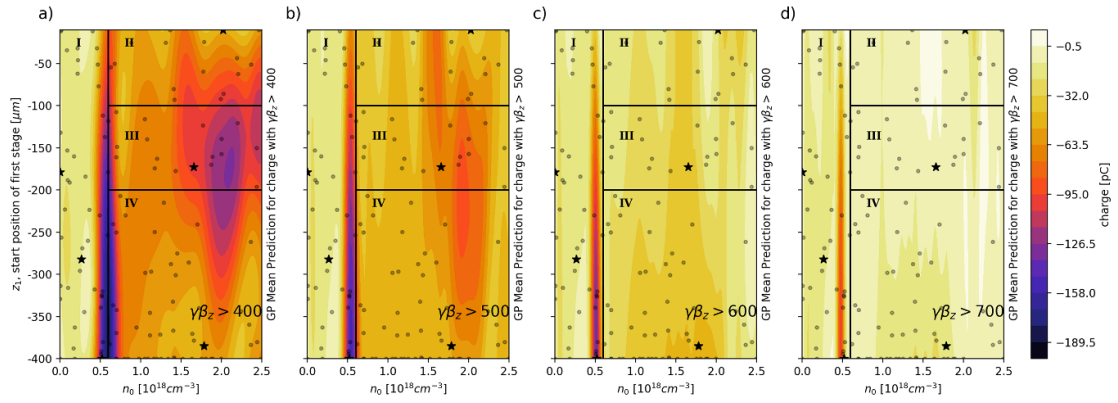


Figure 2. Gaussian process model for the charge at the end of the second stage, or $z = 2.54$ mm. Each subplot shows the charge with momenta above a different momentum cutoff. The scatter plot points show parameters where a simulation was run. Black stars specify the simulations selected as examples for figure 8. Lines divide the regions of the parameter space with labels I, II, III, and IV for the discussion. Region I has the global maximum charge magnitude with 142 pC of charge with at least $\gamma\beta_z > 600$. Region III has significant charge reaching $\gamma\beta_z > 400$ but less charge above the $\gamma\beta_z > 600$ cutoff. Region IV has a local maximum for the magnitude of charge with $\gamma\beta_z > 600$ and > 700 at the bottom of the plot where more charge is accelerated to above those cutoffs.




Figure 3. GP model for maximum energy electron in MeV. Black points show where a simulation was run. White stars show the simulations selected for figure 8.

previously mentioned local maximum in the magnitude of charge with $\gamma\beta_z > 600$ or 700 was located in figure 2.

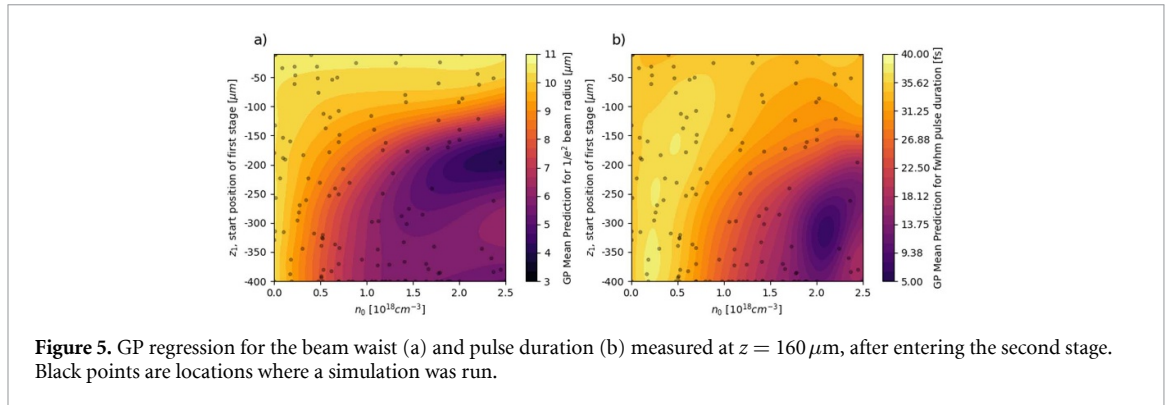
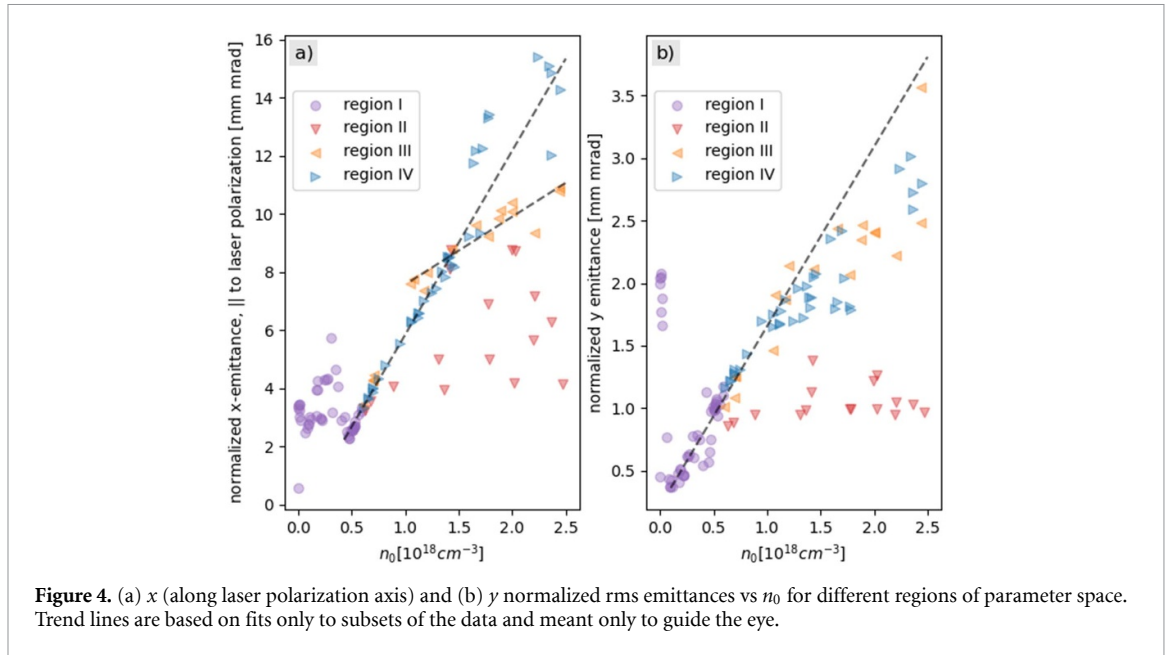
3.4. Emittance results

The normalized transverse emittance is defined [56] as

$$\epsilon_{n,\text{rms},x} = \frac{1}{m_e c} \sqrt{\langle x^2 \rangle \langle p_x^2 \rangle \langle x p_x \rangle^2} \quad (2)$$

where, $\langle \rangle$ denotes the second moment and the x -position and x -momentum x and p_x can be replaced by y and p_y to get $\epsilon_{n,\text{rms},y}$. The normalized emittance was calculated from the electrons in the first bubble at the end of the second stage at $z = 2.54$ mm and is shown in figure 4.

The emittance is generally higher in the x -direction, along the laser polarization, than the y -direction. The x -emittance shows a linear dependence on n_0 for the blue (right facing triangular) and purple (circular) points in figure 4 starting at around $n_0 = 0.48 \times 10^{18} \text{ cm}^{-3}$ with a slope of 6.33 mm mrad per 10^{18} cm^{-3} . The orange (left facing triangular) points, which are for simulations with $-200 \mu\text{m} < z_1$



$< -100 \mu\text{m}$ follow a trend line with a smaller slope of $2.33 \text{ mm mrad per } 10^{18} \text{ cm}^{-3}$. Below $n_0 = 0.48 \times 10^{18} \text{ cm}^{-3}$, the x -emittance is not following the trend line but remains in the range of $2\text{--}6 \text{ mm mrad}$, except for the baseline point where $n_0 = 0$ which indicates that there is no first stage. This point had the lowest x -emittance of 0.58 mm mrad . Additionally, as the first stage is made more thin, the x -emittance becomes less dependent on the density, resulting in the triangular region of the plot below the trend lines for the orange and blue points where the red (down facing triangular) points lie, corresponding to $z_1 > -100 \mu\text{m}$.

In the y -emittance, the linear increase in the emittance starts at a lower value of $n_0 \approx 0.1 \times 10^{18} \text{ cm}^{-3}$. The emittance increases linearly until $n_0 \approx 1 \times 10^{18} \text{ cm}^{-3}$ where points start to fall below the trendline. The trend line in the figure has a slope of $1.44 \text{ mm mrad per } 10^{18} \text{ cm}^{-3}$. There are notable outliers from this trend below $n_0 \approx 0.1 \times 10^{18} \text{ cm}^{-3}$ where the emittance increases from the $n_0 = 0$ value of 0.45 mm mrad to 2 mm mrad before decreasing again to 0.37 mm mrad at the start of the linear increasing trend. Other significant outliers are the red (down facing triangular) points which have a relatively flat dependence on n_0 .

3.5. Pulse propagation results

A GP model for the pulse duration after the end of the first stage and an additional $160 \mu\text{m}$ of propagation in the second stage is shown in figure 5. The additional $160 \mu\text{m}$ of propagation in the second stage contributes little non-linearity to the pulse but enables some divergence. The pulse duration of the laser is self-compressed in the first stage, with the most compression occurring in the thicker and higher density regions of the parameter space. The FWHM pulse duration gets as low as 10 fs at the selected timestep for a compression factor of $3.5 \times$ compared to the original pulse duration. The field energy at this time-step retains $>80\%$ of the original 1 J field energy (the simulation with the most pulse compression had 86% for example) leading to an increase in the peak power.

4. Discussion

The first stage affects the acceleration in the subsequent stage through both injection of electrons into the wake as well as shaping the laser pulse. In order to understand these processes, it is important to understand the relationship between the electron density and the plasma frequency, and how this impacts the plasma wavelength. The electron density controls the plasma wavelength and the density gradients that form the wakefield structure, with higher electron density leading to a shorter plasma wavelength, and steeper density gradients. The plasma wavelength is crucial for the injection and acceleration processes since it controls the distance between the laser and the wake, changing both the dephasing length and the region of the phase space where injection can occur. Additionally, a steeper electron density gradient at higher density increases the rate that the laser pulse is compressed [22] and self-focused. However, too high of a density can lead to effects like beam breakup, which adversely affects the ability of the pulse to drive a wakefield in the subsequent stage. Scanning the density and the length-scale of the first stage therefore enables us to tune the resulting electron beam's properties.

There are multiple injection processes that we would expect to be present in an LWFA with the species and density profile used in these simulations. We expect the dominant injection mechanisms to be ionization injection due to the presence of the high Z argon and density transition injection [30], alternatively referred to as a down-ramp injection in the literature [33, 34], which occurs when the plasma density decreases at the transition between the two stages. In a constant density low Z gas, self-injection from electrons scattering into the bubble from the wake [57] would be the primary injection process. However, for the simulations presented in this article, the density of the helium in the second stage for these laser parameters is below the self-injection threshold given in [58] so we would not expect it to be the dominant mechanism when ionization injection or density transition injection is present.

To ease discussion, we divide the parameter space into different non-overlapping regions labeled I, II, III, and IV where we observed different patterns in both the momentum spectra and the emittance during our exploration of the parameter space. Region I is located on the lower density side of figure 2 with $n_0 < 6 \times 10^{17} \text{ cm}^{-3}$ where we identified the most electrons with $\gamma\beta_z > 600$. In this region, we see a transition between different injection mechanisms depending on the value of n_0 . The remaining three regions have $n_0 > 6 \times 10^{17} \text{ cm}^{-3}$ but are split into different intervals of z_1 since the resulting electron spectra above this density exhibited qualitatively different trends based on the thickness of the first stage. We split the $n_0 > 6 \times 10^{17} \text{ cm}^{-3}$ parameter space into region II with $-100 \mu\text{m} < z_1 < -10 \mu\text{m}$, region IV with $z_1 < -200 \mu\text{m}$ and region III in between II and IV. The regions are annotated in figures 2 and 3. Region I, II, and III had relatively little laser pulse compression, whereas high levels of compression were observed in region IV.

The largest magnitude of charge with $\gamma\beta_z > 600$ occurs within an optimal band of region I near $n_0 \approx 0.5 \times 10^{18} \text{ cm}^{-3}$ which can be seen in figure 2(c). This band has features that we would expect from a density transition process as the change in n_0 affects the density ratio between the two stages which in turn controls the change in λ_p . The magnitude of the λ_p increase affects the location that the electrons are injected in the second stage because when λ_p becomes larger, electrons trapped in the first stage wake are now ahead of the wake in the second stage when the bubble expands. A smaller change in λ_p between stages enables injection closer to the back of the bubble in the second stage which maximizes the energy gain, but a smaller change in λ_p also has a tradeoff in the quantity of charge that can be injected as there is a smaller volume for injection. This explains why as the momentum cutoff is increased in figure 2, the band of n_0 where the most charge is injected narrows and moves left as decreasing n_0 increases λ_p in the first stage. In this lower density region of the parameter space, the laser pulse is not significantly modulated as the rate of self-focusing and pulse compression depends on the electron density.

At higher first stage densities, the electrons are injected into the second stage bubble further ahead of the wake, since λ_p in the first stage is shorter. As a result, even though significant charge is injected the maximum possible energy gain for the individual electrons is lower as the electrons start out partially de-phased. This feature can be seen in figure 2(a) where in region III, there is significant charge in the lower momentum cutoff of $\gamma\beta_z > 400$. However in region III, much less charge compared to the other regions reaches the higher momentum $\gamma\beta_z > 700$ cutoff in figure 2(d). However, as the density of the first stage is further increased, the electrons in the second period of the wakefield in the first stage can be phase-matched to be injected into the back of the first wake of the second stage [59]. We would then expect to see a second optima in the higher momentum cutoffs at this higher density. However, by looking at the $\gamma\beta_z > 700$ plot in figure 2(d) or the slice in figure 6 we see that this second optima only appears in region IV.

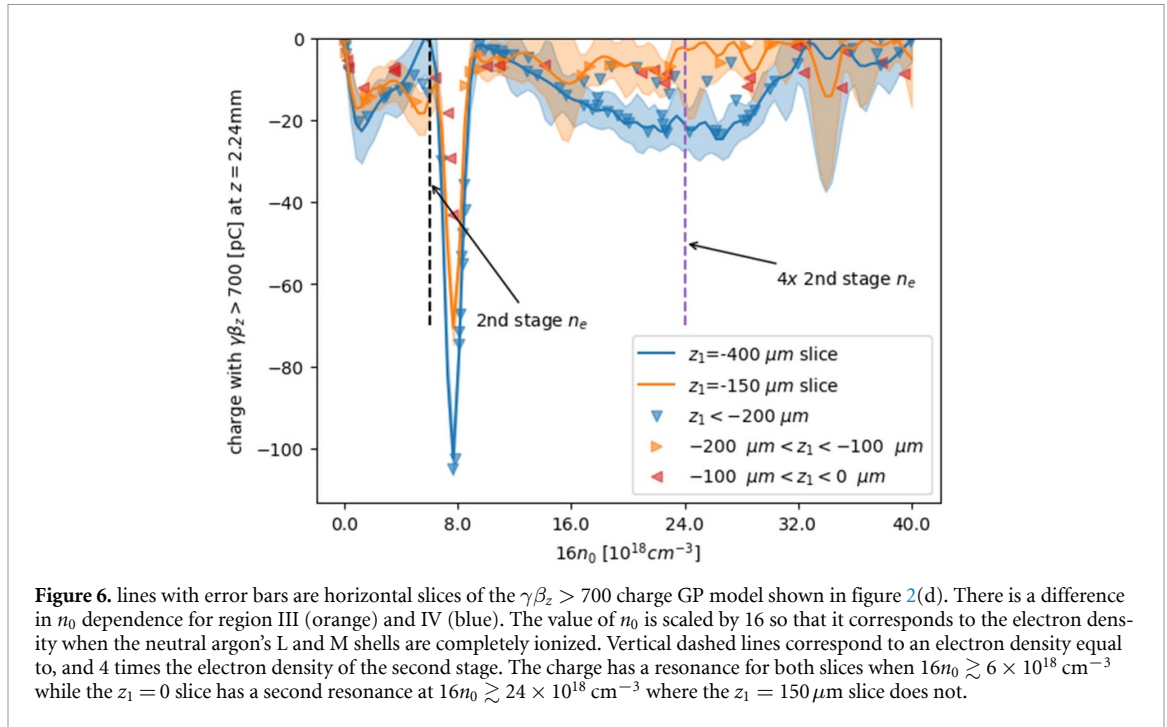
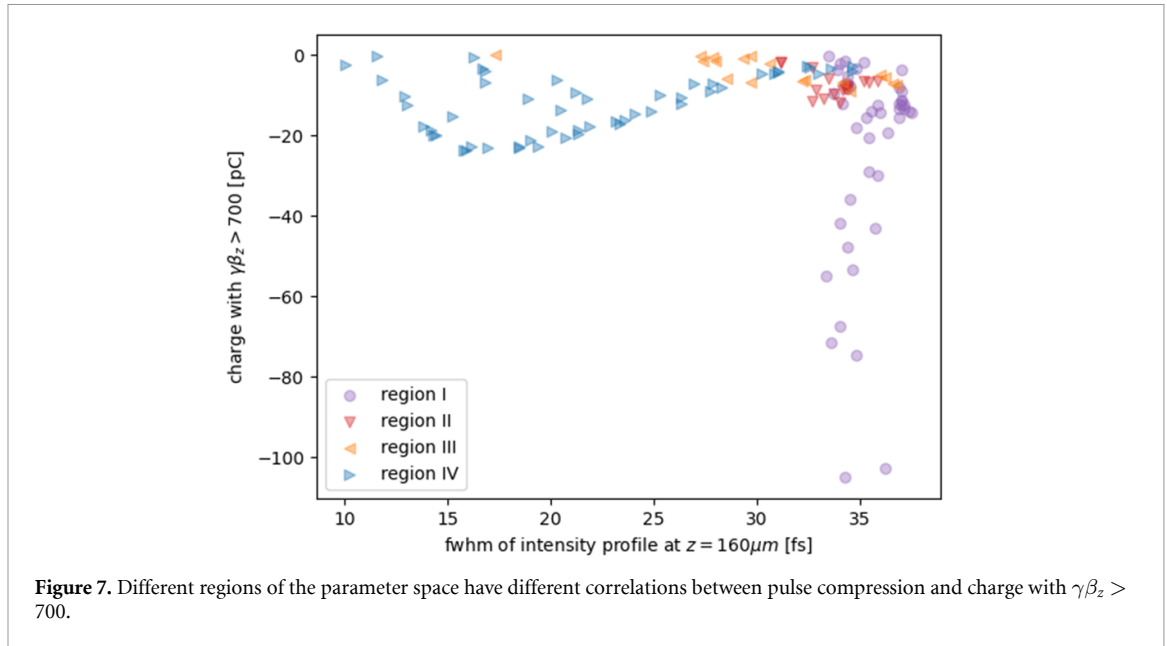


Figure 6. lines with error bars are horizontal slices of the $\gamma\beta_z > 700$ charge GP model shown in figure 2(d). There is a difference in n_0 dependence for region III (orange) and IV (blue). The value of n_0 is scaled by 16 so that it corresponds to the electron density when the neutral argon's L and M shells are completely ionized. Vertical dashed lines correspond to an electron density equal to, and 4 times the electron density of the second stage. The charge has a resonance for both slices when $16n_0 \gtrsim 6 \times 10^{18} \text{ cm}^{-3}$ while the $z_1 = 0$ slice has a second resonance at $16n_0 \gtrsim 24 \times 10^{18} \text{ cm}^{-3}$ where the $z_1 = 150 \mu\text{m}$ slice does not.

In figures 2(c) and (d), corresponding to momentum cutoffs $\gamma\beta_z > 600$ and $\gamma\beta_z > 700$, there is a local optimum in the quantity of charge in region IV that does not occur at the same n_0 value for region III. This can be seen in figure 6 where the $z_1 = -400 \mu\text{m}$ slice in blue has a local optima around $16n_0 \gtrsim 24 \times 10^{18} \text{ cm}^{-3}$ while $z_1 = 150 \mu\text{m}$ slice does not. Figure 6 shows that this optima occurs when the electron density at the transition is $\sim 4\times$ the electron density of the second stage. An explanation for this difference may be inferred by looking at figure 7, where there is a different correlation between the pulse compression in the first stage and the resulting accelerated charge for different regions of the parameter space. In particular, region IV has an optimum for the charge at a FWHM of 15 fs while this trend does not appear in points from region III, despite both results spanning a similar n_0 range. The longer first stage in region IV provides more length for the pulse to compress such that the pulse is shorter during the transition between the two stages. There are two primary means that a shortening pulse duration can enhance the injection. First, a longer propagation in the first stage enables the pulse to get closer to, or satisfy, the $c\tau < \lambda_p/2$ requirement for efficiently driving the wake, which is a stricter condition as the density is increased (τ is the FWHM pulse duration). Another benefit for injection is that increasing a_0 leads to an increase in the bubble radius due to relativistic corrections to the plasma frequency [14, 22]. The increase in the bubble radius from the increasing a_0 can have a similar effect to a density transition on the injection.

In region II, the plasma wavelength, λ_p is between 5 and 11 μm at the end of the first stage, depending on n_0 . Meanwhile the thickness of the region of z with argon is between 10 μm and 100 μm . As the density of the argon must increase from 0 to n_0 over that length, the up-ramp's length scale becomes comparable to λ_p . This has implications for the laser pulse, which is relatively unmodified by the first stage, as can be seen in figure 5. Additionally, the density transition injection process is impacted, as the bubble structure has less space to form in the first stage.

The distinct physical mechanisms underlying the behavior in each region can be directly observed in the electron density at several timesteps and in the resulting momentum spectra of representative simulation examples as shown in figure 8. The electron density is plotted to show the wake structure at three times: the accelerating bunch at the transition between the two stages, after the wake is fully formed in the accelerating stage, and at the end of the accelerator length. Meanwhile, the electron momentum spectra show the impact of the injection mechanism on the momentum spread and quantity of charge. First, it is useful to consider the case of a homogeneous accelerator, i.e. a single stage without an injection stage which is shown in figure 8(a). The remaining cases can then be benchmarked against this initial case. In the absence of a first stage, equivalent to when n_0 is set to 0, charge is self-injected while the laser propagates in the helium accelerating stage. This configuration results in a relatively low charge but high maximum $\gamma\beta_z$, as the minimal beam-loading enables a stronger accelerating field. This is a baseline case which we can compare the other examples to.



We can also plot the case from our simulations, also from region I, which produced the highest energy electrons in figure 8(b). This example shows similar behavior to the case with no first stage but with more charge overall and a slightly higher maximum $\gamma\beta_z$. This simulation had a n_0 value that results in a lower electron density on the first stage side of the transition than in the second stage. Density transition injection can therefore be ruled out since the transition is from a larger λ_p to a smaller λ_p . The higher charge compared to the baseline without a first stage is because the charge is being injected through ionization injection. In ionization injection, electrons with higher appearance intensity are ionized in the peak of the laser pulse rather than the leading edge which enables them to fall behind the pulse and into the wake [31, 32]. We can confirm that this is what occurs in this example by tracking the particles as shown in the supplemental. As the electrons are injected in the first stage, some of the electrons get more acceleration length and reach higher energies compared to the baseline case, where the electrons are injected at some point in the second stage.

Using the same approach we can examine representative cases in each of the regions with a transition from a higher to a lower electron density to validate our hypotheses for injection and acceleration in each case. Figure 8(c) shows a representative simulation in region I where we begin to observe density transition injection. We chose to display this example because it had the most charge with $\gamma\beta_z > 600$, which was the condition that the algorithm we used was optimizing. When electrons are injected into the same region at the back of the wake, as can be seen in the second time step's electron density, and the injection is localized to the location of the transition the resulting energy spread is mono-energetic. We can see that this is what happens in this example as the momentum spread is narrow even as the total charge is quite high compared to the other examples.

In the region II example in figure 8(d), the first stage is so thin that a bubble does not form in the first stage like it would in the other examples, but there is still a density spike that is able to be trapped in the second stage. In the density plot in the second column, the injected electrons extends further from the wake of the bubble which leads to a larger energy spread and relatively low maximum energy. However, the beam remains high charge and has less transverse spread than other cases at a similar n_0 below it.

Finally, in regions III and IV the examples in figures 8(e) and (f) have a similar n_0 but different z_1 . Despite the density ratio being similar at the transition, the region IV example shows two separate bunches while in the region III example only a single bunch is observed. The presence of a double bunch in the region IV example suggests the effect discussed previously where a bunch is injected from the second period in the first stage to the bubble in the second stage occurs. Despite the similar value of n_0 , at the transition, the higher intensity in the region IV case, due to the larger pulse compression and self-focusing, visibly modulates the plasma waves behind the laser differently between the two cases. The end result is that while there is an overall lower charge in the region IV case, the double bunch structure creates a second peak in the momentum spectra that leads to a higher maximum momentum from the charge that is injected closer to the back of the wake.

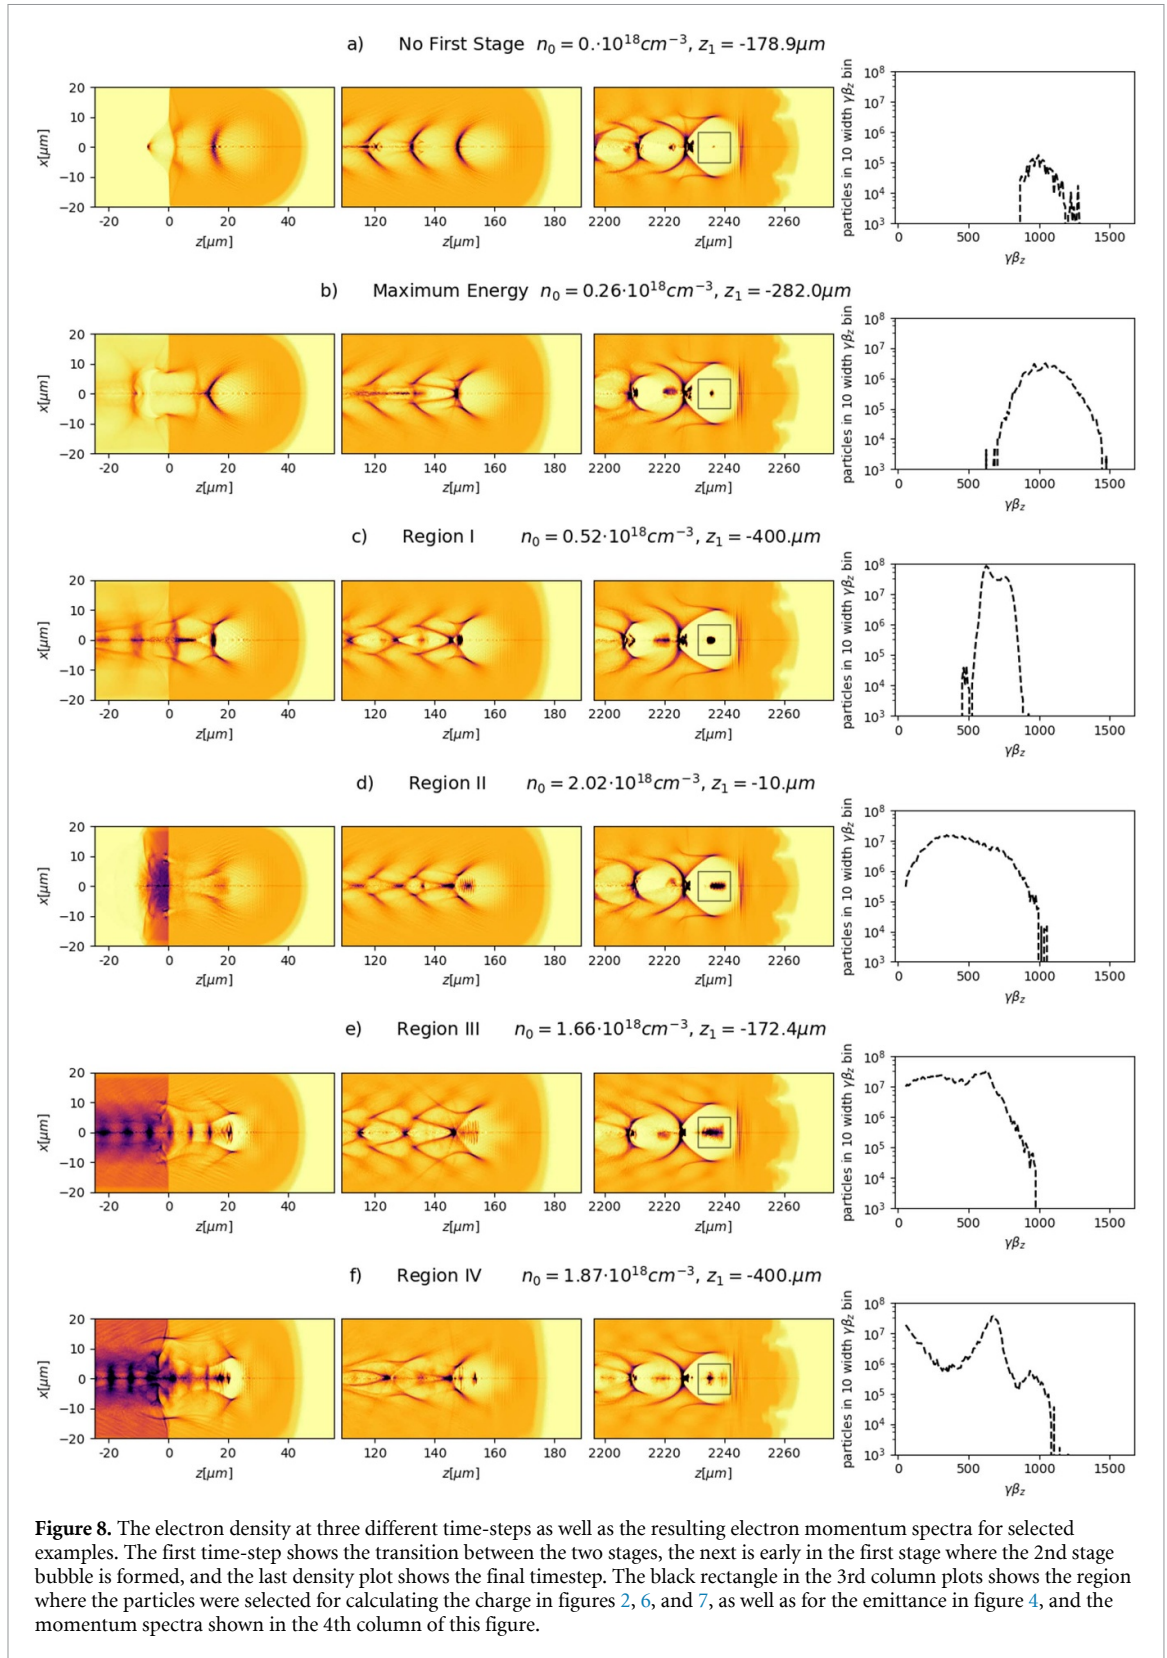


Figure 8. The electron density at three different time-steps as well as the resulting electron momentum spectra for selected examples. The first time-step shows the transition between the two stages, the next is early in the first stage where the 2nd stage bubble is formed, and the last density plot shows the final timestep. The black rectangle in the 3rd column plots shows the region where the particles were selected for calculating the charge in figures 2, 6, and 7, as well as for the emittance in figure 4, and the momentum spectra shown in the 4th column of this figure.

4.1. Discussion of emittance trends

By plotting a scatter plot of the emittance vs n_0 separated by region as in figure 4, we can correlate the various injection mechanisms observed in each region to the emittance.

The overall higher emittance in the direction of the laser polarization observed for all regions in figure 4 is expected. As discussed in [32], ionization injection, which would be expected in the argon first stage but not in the helium second stage, typically leads to higher divergence in the direction of the laser polarization as the electrons are born in the laser field and interact with it before falling back to

be trapped in the wake. Interactions with the laser field result in increased transverse momentum and position spread and would be expected to impact the emittance as well.

In figure 4, the emittance in both x and y increases with n_0 from the baseline $n_0 = 0$ case to a local maximum at the lowest but non-zero n_0 value. Subsequently, the x and y emittance have a different pattern where the y -emittance has a linearly increasing trend, while the x -emittance has a second peak before it also follows a linear trend that begins at a higher value of n_0 . We hypothesize that the initial growth in the emittance with n_0 between the helium only case and non-zero n_0 is due to a switch to ionization injection from self injection. The x -emittance is significantly higher than the y -emittance for $n_0 > 0$, as opposed to the pure helium accelerating stage where the emittance in both directions are comparable. Because the ionization injection process leads to interaction with the laser field, the x -emittance is more sensitive to the details of where the electrons that end up injected are distributed. This helps to explain the second peak in the x -emittance as the first stage bubble radius changes with n_0 , which changes the interaction between the electrons at the front of the bubble and the laser modifying the trajectories of the electrons. Subsequently, the x -emittance decreases until a linear increasing trend begins. This linear trend in the x -emittance starts at the same density that we see the optimal band for the charge in figure 2, suggesting that the trend is due to the density transition injection process becoming the dominant injection mechanism. Density transition injection would have increasing emittance with the value of n_0 because a larger density ratio between the two stages leads to a larger bubble expansion at the transition, which enables electrons located in a larger transverse area (as well as longitudinal) to be injected. This may explain the linear increase in the emittance with n_0 in region I, III, and IV.

The emittance in region II is well below the trends for region III and IV, with the y -emittance especially showing a clear separation between the emittance trend in region II compared to the other regions. More qualitatively, there is also a clearly narrow transverse extent of the bunch in the region II example of figure 8 than the subsequent examples.

5. Conclusion

We have shown how the density and thickness of an injector stage of a realistic two-stage LWFA controls the resulting electron beam properties via hybridization of density transition and ionization injection mechanisms. Use of numerical optimization schemes and GPs enables us to identify distinct physical regimes of LWFA in a complex parameter space. By scanning just two parameters we identify regions of the parameter space enabling desired properties of the electron beam like the maximum energy, charge, or emittance to be tuned. In region I, when the first stage is at a lower density than the second, the highest energy electrons were observed. When the first stage density is increased in region I, we observe an ideal range of densities for density transition injection, where the low emittance, high charge, and lowest momentum spread bunches are observed. As the density is further increased, we observe that for thin first stages in region II, we see broad energy spread low emittance electrons. At the same densities, as the length of the first stage is increased, the emittance increases but the charge also increases in region III. Eventually, when the density and length of the first stage is long enough, we see a region with significant pulse compression in region IV where a double bunch structure forms which, while having less charge overall than region III, enhances the amount of charge on the high momentum end of the spectrum.

Running two simulations in parallel, each on a NVIDIA V100 GPU, the full parameter scan required less than 6 d, demonstrating that formerly intractable LWFA parameter scans may now be completed with practical amounts of time and resources. As laser systems increase in repetition rates, numerical methods will be needed that are similarly repetition rated.

5.1. Future work on multi-stage optimization

For simplicity, this work restricts the number of parameters to just two: the starting position and density of the first stage. It is exciting to consider the possibilities of extending the technique to include other impactful parameters. For instance, a longer transition between two stages has been shown to reduce the emittance at a cost of reducing the quantity of charge [29], adding the transition length as a parameter could further improve LWFA performance. Ionization injection can be tuned by varying different gas mix ratios. In general, for higher dimensional parameter spaces accurately identifying regions of interest and local optima become more time consuming as more scans are needed compared to the two dimensional case studied here. Despite that, the GP regression and Bayesian optimization method used here is well suited for optimizing a multidimensional parameter space that varies all of these different parameters while minimizing the number of required simulations and thus the required time.

There is also significant interest in scaling LWFA to the multi-PW regime to reach energies at or exceeding 10's of GeV. Future work could look at optimizing the density transition mechanism for controlling beam quality with a multi-Joule beam at a lower second stage density. The lower density would be required to reduce self focusing and would also have the benefit of increasing the dephasing length for higher energy gain. Understanding how to effectively surpass the dephasing limit would enable substantially higher electron energies, and would likely require careful examination of the parameter space for how to structure the laser or plasma density to accomplish such a feat.

5.2. Scaling pulse compression and multi-stage LWFA to higher energy lasers

It is of note that in these simulations significant pulse compression can be achieved in the first stage in region IV with FWHM pulse durations as low as 10 fs compared to the original 35 fs of the input pulse. Achieving pulse durations below 20 fs with conventional Ti-Sapphire based chirped pulse amplification [60] lasers at Joule level energies is extremely challenging. Post-compression after amplification uses non-linear effects to achieve high power few cycle laser pulses, such as thin film compression [61]. Scaling post-compression after amplification techniques to high energy pulses [20] while maintaining throughput efficiency and beam quality is a grand challenge problem for laser technology [62]. For high peak power laser pulses, non-uniformity in the near field makes it difficult to use most optics to pulse compress as the beam must stay under damage thresholds and non-uniformity results in varying spectral broadening across the pulse [19]. Demonstrating the usability of such a pulse after using non-linear effects to broaden the laser spectrum is also challenging since spatial non-uniformity [63] and spatio-temporal couplings [64] can be introduced. A setup like the one simulated here could demonstrate both a significant factor of pulse compression through interactions in the far field while avoiding damage thresholds by using plasma, and subsequently using acceleration of electrons in the second stage to demonstrate the usability of the compressed pulse.

Recently, laser facilities with petawatt- or even multi-petawatt- powers are becoming available [65–69]. It would be a challenge to compress such high energy pulses using the density gradients from LWFA. As the bubble created by the pulse has similar spatial dimensions both longitudinally and transversely, increasing the density to make a small plasma wavelength that can temporally compress the pulse also leads the pulse to self focus. At higher pulse energy, this self focusing also will increase the a_0 , and the resulting relativistic correction to the plasma frequency will elongate the plasma wavelength. The result is an elongation of the bubble longitudinally that reduces the ratio of temporal compression to self focusing. The density would then have to be lowered to reduce self focusing which would necessitate a larger propagation length to do much compression. As the propagation length increases, it becomes more difficult to maintain the beam quality as nonlinearities accumulate. As a result, more effort is required to understand how to implement a technique to perform meaningful pulse compression on multi-Joule ultra-fast lasers.

Data availability statement

The data cannot be made publicly available upon publication because the cost of preparing, depositing and hosting the data would be prohibitive within the terms of this research project. The data that support the findings of this study are available upon reasonable request from the authors.

Injection analysis available at <https://doi.org/10.1088/1367-2630/ae23f1/data1>.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy under Grant No. DE-SC0023504. This work was supported by the NSF under the following awards; STROBE: A National Science Foundation Science & Technology Center under Grant Nos. DMR-1548924, PHY-1753165, and PHY-2308982.

This research used the open-source particle-in-cell code WarpX <https://github.com/ECP-WarpX/WarpX>. Primary WarpX contributors are with LBNL, LLNL, CEA-LIDYL, SLAC, DESY, CERN and TAE Technologies. We acknowledge all WarpX contributors.

This work utilized the infrastructure for high-performance and high-throughput computing, research data storage and analysis, and scientific software tool integration built, operated, and updated by the Research Cyberinfrastructure Center (RCIC) at the University of California, Irvine (UCI). The RCIC provides cluster-based systems, application software, and scalable storage to directly support the UCI research community. <https://rcic.uci.edu>.

ORCID iDs

Christopher Gardner  0000-0002-7983-8265

Franklin Dollar  0000-0003-3346-5763

References

- [1] Tajima T and Dawson J M 1979 Laser electron accelerator *Phys. Rev. Lett.* **43** 267–70
- [2] Kneip S *et al* 2009 Near-GeV acceleration of electrons by a nonlinear plasma wave driven by a self-guided laser pulse *Phys. Rev. Lett.* **103** 35002
- [3] Leemans W P, Nagler B, Gonsalves A J, Tóth C, Nakamura K, Geddes C G, Esarey E, Schroeder C B and Hooker S M 2006 GeV electron beams from a centimetre-scale accelerator *Nat. Phys.* **2** 696–9
- [4] Lundh O *et al* 2011 Few femtosecond, few kiloampere electron bunch produced by a laser-plasma accelerator *Nat. Phys.* **7** 219–22
- [5] He Z H, Thomas A G, Beaurepaire B, Nees J A, Hou B, Malka V, Krushelnick K and Faure J 2013 Electron diffraction using ultra-fast electron bunches from a laser-wakefield accelerator at kHz repetition rate *Appl. Phys. Lett.* **102** 064104
- [6] Rousse A *et al* 2004 Production of a keV x-ray beam from synchrotron radiation in relativistic laser-plasma interaction *Phys. Rev. Lett.* **93** 135005
- [7] Kneip S *et al* 2010 Bright spatially coherent synchrotron x-rays from a table-top source *Nat. Phys.* **6** 980–3
- [8] Albert F and Thomas A G 2016 Applications of laser wakefield accelerator-based light sources *Plasma Phys. Control. Fusion* **58** 103001
- [9] Geddes C G, Toth C, Tilborg J V, Esarey E, Schroeder C B, Bruhwiler D, Nieter C, Cary J and Leemans W P 2004 High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding *Nature* **431** 538–41
- [10] Faure J, Glinec Y, Pukhov A, Kiselev S, Gordienko S, Lefebvre E, Rousseau J P, Burgy F and Malka V 2004 A laser-plasma accelerator producing monoenergetic electron beams *Nature* **431** 541–4
- [11] Mangles S P *et al* 2004 Monoenergetic beams of relativistic electrons from intense laser-plasma interactions *Nature* **431** 535–8
- [12] Pukhov A and Meyer-ter-Vehn J 2002 Laser wake field acceleration: the highly non-linear broken-wave regime *Appl. Phys. B* **74** 355–61
- [13] Lu W, Huang C, Zhou M, Tzoufras M, Tsung F S, Mori W B and Katsouleas T 2006 A nonlinear theory for multidimensional relativistic plasma wave wakefields *Phys. Plasmas* **13** 056709
- [14] Lu W, Tzoufras M, Joshi C, Tsung F S, Mori W B, Vieira J, Fonseca R A and Silva L O 2007 Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime *Phys. Rev. Spec. Top.-Accel. Beams* **10** 061301
- [15] Salehi F, Goers A J, Hine G A, Feder L, Kuk D, Miao B, Woodbury D, Kim K Y and Milchberg H M 2017 MeV electron acceleration at 1 kHz with <10 mJ laser pulses *Opt. Lett.* **42** 215–8
- [16] Guénot D *et al* 2017 Relativistic electron beams driven by kHz single-cycle light pulses *Nat. Photon.* **11** 293–6
- [17] Beaurepaire B, Lifschitz A and Faure J 2014 Electron acceleration in sub-relativistic wakefields driven by few-cycle laser pulses *New J. Phys.* **16** 023023
- [18] Lifschitz A F and Malka V 2012 Optical phase effects in electron wakefield acceleration using few-cycle laser pulses *New J. Phys.* **14** 053045
- [19] Khazanov E A, Mironov S Y and Mourou G 2019 Nonlinear compression of high-power laser pulses: compression after compressor approach *Phys.-Usp.* **62** 1096–124
- [20] Farinella D M *et al* 2019 Focusability of laser pulses at petawatt transport intensities in thin-film compression *J. Opt. Soc. Am. B* **36** A28
- [21] Mori W B 1997 The physics of the nonlinear optics of plasmas at relativistic intensities for short-pulse lasers *IEEE J. Quantum Electron.* **33** 1942–53
- [22] Schreiber J, Bellei C, Mangles S P D, Kamperidis C, Kneip S, Nagel S R, Palmer C A J, Rajeev P P, Streeter M J V and Najmudin Z 2010 Complete temporal characterization of asymmetric pulse compression in a laser wakefield *Phys. Rev. Lett.* **105** 235003
- [23] Pipahl A, Anashkina E A, Toncian M, Toncian T, Skobelev S A, Bashinov A V, Gonoskov A A, Willi O and Kim A V 2013 High-intensity few-cycle laser-pulse generation by the plasma-wakefield self-compression effect *Phys. Rev. E* **87** 033104
- [24] Malka V 2012 Laser plasma accelerators *Phys. Plasmas* **19** 055501
- [25] Geddes C G R, Nakamura K, Plateau G R, Toth C, Cormier-Michel E, Esarey E, Schroeder C B, Cary J R and Leemans W P 2008 Plasma-density-gradient injection of low absolute-momentum-spread electron bunches *Phys. Rev. Lett.* **100** 215004
- [26] Massimo F, Lifschitz A F, Thaury C and Malka V 2017 Numerical studies of density transition injection in laser wakefield acceleration *Plasma Phys. Control. Fusion* **59** 085004
- [27] Buck A *et al* 2013 Shock-front injector for high-quality laser-plasma acceleration *Phys. Rev. Lett.* **110** 185006
- [28] Schmid K, Buck A, Sears C M S, Mikhailova J M, Tautz R, Herrmann D, Geissler M, Krausz F and Veisz L 2010 Density-transition based electron injector for laser driven wakefield accelerators *Phys. Rev. Spec. Top.-Accel. Beams* **13** 091301
- [29] Samant S A, Upadhyay A K and Krishnagopal S 2014 High brightness electron beams from density transition laser wakefield acceleration for short-wavelength free-electron lasers *Plasma Phys. Control. Fusion* **56** 095003
- [30] Suk H, Barov N, Rosenzweig J B and Esarey E 2001 Plasma electron trapping and acceleration in a plasma wake field using a density transition *Phys. Rev. Lett.* **86** 1011–4
- [31] McGuffey C *et al* 2010 Ionization induced trapping in a laser wakefield accelerator *Phys. Rev. Lett.* **104** 025004
- [32] Pak A, Marsh K A, Martins S F, Lu W, Mori W B and Joshi C 2010 Injection and trapping of tunnel-ionized electrons into laser-produced wakes *Phys. Rev. Lett.* **104** 025003
- [33] Hansson M, Aurand B, Davoine X, Ekerfelt H, Svensson K, Persson A, Wahlström C and Lundh O 2015 Down-ramp injection and independently controlled acceleration of electrons in a tailored laser wakefield accelerator *Phys. Rev. Spec. Top.-Accel. Beams* **18** 71303
- [34] Gonsalves A J *et al* 2011 Tunable laser plasma accelerator based on longitudinal density tailoring *Nat. Phys.* **7** 862–6
- [35] Reitsma A J W, Leemans W P, Esarey E, Schroeder C B, Kamp L P J and Schep T J 2002 Simulation of electron postacceleration in a two-stage laser wakefield accelerator *Phys. Rev. Spec. Top.-Accel. Beams* **5** 051301
- [36] Pollock B B *et al* 2011 Demonstration of a narrow energy spread, 0.5 GeV electron beam from a two-stage laser wakefield accelerator *Phys. Rev. Lett.* **107** 045001

- [37] Chien T, Chang C, Lee C, Lin J, Wang J and Chen S 2005 Spatially localized self-injection of electrons in a self-modulated laser-wakefield accelerator by using a laser-induced transient density ramp *Phys. Rev. Lett.* **94** 115003
- [38] Irshad F, Karsch S and Döpp A 2023 Multi-objective and multi-fidelity bayesian optimization of laser-plasma acceleration *Phys. Rev. Res.* **5** 13063
- [39] Pousa A F, Jalias S, Kirchen M, de la Ossa A, Thévenet M, Hudson S, Larson J, Huebl A, Vay J and Lehe R 2023 Bayesian optimization of laser-plasma accelerators assisted by reduced physical models *Phys. Rev. Accel. Beams* **26** 84601
- [40] Irshad F, Eberle C, Foerster F, Grafenstein K, Haberstroh F, Travac E, Weisse N, Karsch S and Döpp A 2024 Pareto optimization and tuning of a laser wakefield accelerator *Phys. Rev. Lett.* **133** 85001
- [41] Lindstrom C A 2021 Staging of plasma-wakefield accelerators *Phys. Rev. Accel. Beams* **24** 014801
- [42] Steinke S *et al* 2016 Multistage coupling of independent laser-plasma accelerators *Nature* **530** 190–3
- [43] Fedeli L *et al* 2022 Pushing the frontier in the design of laser-based electron accelerators with groundbreaking mesh-refined particle-in-cell simulations on exascale-class supercomputers *Int. Conf. for High Performance Computing, Networking, Storage and Analysis*, SC vol 2022–November (IEEE Computer Society)
- [44] Vay J *et al* 2021 Modeling of a chain of three plasma accelerator stages with the WarpX electromagnetic PIC code on GPUs *Phys. Plasmas* **28** 023105
- [45] Lehe R, Lifschitz A, Thaury C, Malka V and Davoine X 2013 Numerical growth of emittance in simulations of laser-wakefield acceleration *Phys. Rev. Spec. Top.-Accel. Beams* **16** 021301
- [46] Kirchen M, Lehe R, Jalias S, Shapoval O, Vay J and Maier A R 2020 Scalable spectral solver in Galilean coordinates for eliminating the numerical Cherenkov instability in particle-in-cell simulations of streaming plasmas *Phys. Rev. E* **102** 13202
- [47] Ammosov M V, Delone N B and Krainov V P 1986 Tunnel ionization of complex atoms and atomic ions in electromagnetic field *Proc. SPIE* **0664** 138
- [48] Mangles S P 2014 An overview of recent progress in laser wakefield acceleration experiments *CAS-CERN Accelerator School: Plasma Wake Acceleration 2014, Proc.* (CERN) pp 289–300
- [49] Vargas M *et al* 2014 Improvements to laser wakefield accelerated electron beam stability, divergence and energy spread using three-dimensional printed two-stage gas cell targets *Appl. Phys. Lett.* **104** 174103
- [50] He Z H, Nees J A, Hou B, Krushelnick K and Thomas A G 2014 Enhancement of plasma wakefield generation and self-compression of femtosecond laser pulses by ionization gradients *Plasma Phys. Control. Fusion* **56** 084010
- [51] Augst S, Strickland D, Meyerhofer D D, Chin S L and Eberly J H 1989 Tunneling ionization of noble gases in a high-intensity laser field *Phys. Rev. Lett.* **63** 2212–5
- [52] Rasmussen C E and Williams C K I 2005 *Gaussian Processes for Machine Learning* (MIT Press)
- [53] Boltz T, Martinez J L, Xu C, Baker K R L, Roussel R, Ratner D, Mustapha B and Edelen A L 2024 More sample-efficient tuning of particle accelerators with Bayesian optimization and prior mean models (arXiv:2403.03225)
- [54] Hudson S, Larson J, Navarro J L and Wild S M 2022 Libensemble: a library to coordinate the concurrent evaluation of dynamic ensembles of calculations *IEEE Trans. Parallel Distrib. Syst.* **33** 977–88
- [55] Pedregosa F *et al* 2011 Scikit-learn: machine learning in python *J. Mach. Learn. Res.* **12** 2825–30 (available at: <http://jmlr.org/papers/v12/pedregosa11a.html>)
- [56] Floettmann K 2003 Some basic features of the beam emittance *Phys. Rev. Spec. Top.-Accel. Beams* **6** 80–86
- [57] Bulanov S V, Pegoraro F, Pukhov A M and Sakharov A S 1997 Transverse-wake wave breaking *Phys. Rev. Lett.* **78** 4205–8
- [58] Mangles S P D, Genoud G, Bloom M S, Burza M, Najmudin Z, Persson A, Svensson K, Thomas A G R and Wahlström C 2012 Self-injection threshold in self-guided laser wakefield accelerators *Phys. Rev. Spec. Top.-Accel. Beams* **15** 011302
- [59] Deng A H *et al* 2012 Control of electron-seeding phase in a cascaded laser wakefield accelerator *Phys. Plasmas* **19** 023105
- [60] Strickland D and Mourou G 1985 Compression of amplified chirped optical pulses *Opt. Commun.* **55** 447–9
- [61] Mironov S Y, Ginzburg V N, Yakovlev I V, Kochetkov A A, Shaykin A A, Khazanov E A and Mourou G A 2017 Using self-phase modulation for temporal compression of intense femtosecond laser pulses *Quantum Electron.* **47** 614–9
- [62] Kling M F *et al* 2025 Roadmap on basic research needs for laser technology *J. Opt.* **27** 013002
- [63] Mak A A, Soms L N, Fromzel' V A and Iashin V E 1990 *Nd-Glass Lasers* (Moscow Izdatel Nauka)
- [64] Akturk S, Gu X, Bowlan P and Trebino R 2010 Spatio-temporal couplings in ultrashort laser pulses *J. Opt.* **12** 093001
- [65] Yoon J W, Kim Y G, Choi I W, Sung J H, Lee H W, Lee S K and Nam C H 2021 Realization of laser intensity over 10^{23} W/cm² *Optica* **8** 630
- [66] Papadopoulos D N *et al* 2016 The Apollon 10 PW laser: experimental and theoretical investigation of the temporal characteristics *High Power Laser Sci. Eng.* **4** e34
- [67] Nees J *et al* 2021 Zettawatt equivalent ultrashort pulse laser system: an NSF mid-scale facility for laser-driven science in the QED regime 2021 *Conf. on Lasers and Electro-Optics, CLEO 2021 - Proc.* (Institute of Electrical and Electronics Engineers Inc.) (https://doi.org/10.1364/cleo_at.2021.jtu3a.10)
- [68] Doria D, Cernaianu M O, Ghenuche P, Stutman D, Tanaka K A, Ticos C and Ur C A 2020 Overview of ELI-NP status and laser commissioning experiments with 1 PW and 10 PW class-lasers *J. Instrum.* **15** C09053
- [69] Bromage J *et al* 2021 MTW-OPAL: a technology development platform for ultra-intense optical parametric chirped-pulse amplification systems *High Power Laser Sci. Eng.* **9** e45