Numerical analysis of the plasma density profile of a laser wakefield accelerator using multi-objective Bayesian optimization

Semion Tchetovsky, Igor A Andriyash, Jérôme Faure

PAPER • OPEN ACCESS

Numerical analysis of the plasma density profile of a laser wakefield accelerator using multi-objective Bayesian optimization

To cite this article: Semion Tchetovsky *et al* 2026 *Mach. Learn.: Sci. Technol.* **7** 025049

IOP Publishing Mach. Learn.: Sci. Technol. 7 (2026) 025049 https://doi.org/10.1088/2632-2153/ae5146

MACHINE LEARNING
Science and Technology

PAPER

Numerical analysis of the plasma density profile of a laser wakefield accelerator using multi-objective Bayesian optimization

Semion Tchetovsky*, Igor A Andriyash and Jérôme Faure

Laboratoire d'Optique Appliquée (LOA), CNRS, École Polytechnique, ENSTA, Institut Polytechnique de Paris, Palaiseau, France
* Author to whom any correspondence should be addressed.

E-mail: semion.tchetovsky@polytechnique.edu

Keywords: laser wakefield acceleration, multi-objective Bayesian optimization, PIC simulations, beam divergence

OPEN ACCESS

RECEIVED
18 December 2025

REVISED
16 February 2026

ACCEPTED FOR PUBLICATION
12 March 2026

PUBLISHED
7 April 2026

Abstract

Laser wakefield acceleration (LWFA) provides a compact pathway to generate high-energy electron beams, but its performance depends on a large set of coupled, nonlinear parameters. Machine-learning approaches, particularly Bayesian optimization (BO), have become valuable for efficiently exploring this space. In this study, we examine how the longitudinal plasma density profile affects the properties of LWFA electron beams using high-fidelity quasi-3D particle-in-cell simulations guided by BO. We focus specifically on the influence of the injection gradient and the plasma density down-ramp length. This work establishes a systematic methodology for multi-parameter and multi-objectives optimization of LWFAs and offers design guidance for the upcoming laser-plasma acceleration center installation. The framework is general and adaptable to other accelerator configurations and optimization goals.

1. Introduction

Laser wakefield acceleration (LWFA), first proposed by Tajima and Dawson in 1979 [1], represents a groundbreaking approach to charged-particle acceleration. In their seminal work, they demonstrated through theoretical analysis and simulations that an ultraintense laser pulse could excite large-amplitude plasma waves via the ponderomotive force. The resulting plasma wakefields are capable of sustaining electric fields of the order of hundreds of gigavolts per meter, enabling the acceleration of trapped particles to ultra-relativistic energies over millimeter to centimeter-scale distances [2, 3]. The key advantage of this mechanism lies in its compactness: with accelerating gradients up to four orders of magnitude higher than those achievable in conventional radiofrequency linear accelerators, LWFA opens the path toward table-top accelerators and significantly reduced infrastructure costs.

Over the past two decades, tremendous progress has been achieved in the experimental realization and optimization of LWFAs. The early 2000s marked a milestone with the first demonstrations of quasi-monoenergetic electron beams [4–6], establishing the potential of LWFAs to produce high-quality beams suitable for applications. Subsequent developments have focused on increasing beam energy, improving stability, and enhancing reproducibility. High-energy systems have achieved multi-GeV beams [7–13], and later work has demonstrated significant improvements in the stability of such high-energy operation [14]. In parallel, high-repetition-rate systems have been developed, placing emphasis on operational robustness and achieving similarly high levels of beam stability [15, 16].

With these advances, LWFA technology is reaching a level of maturity that enables the design of next-generation facilities. At the Laboratoire d'Optique Appliquée (LOA), the *laser-plasma acceleration center* (LAPLACE) [17] project aims to deliver high-quality electron beams with energies ranging from 50 to 150 MeV at repetition rates up to 100 Hz, targeting applications that require both performance and operational reliability.

Despite these achievements, one of the primary challenges in the field remains the vastness and complexity of the parameter space that governs LWFA performance. Numerous interdependent parameters—such as laser pulse and gas target characteristics—make manual optimization impractical. Traditional

scanning methods are prohibitively time-consuming and often fail to capture nonlinear correlations between parameters. Consequently, the LWFA community has increasingly turned to advanced data-driven techniques [18], particularly machine learning (ML), to accelerate experimental optimization and interpret multidimensional parameter dependencies. Among these methods, Bayesian optimization (BO) has emerged as a particularly promising approach [19–21]. BO is a sequential, probabilistic technique based on Bayesian inference, designed to optimize expensive and noisy black-box functions using a limited number of evaluations [22, 23]. Its ability to model uncertainty and efficiently balance exploration and exploitation makes it highly suitable for both experimental and numerical plasma accelerator studies, including particle-in-cell (PIC) simulations where computational cost is substantial.

The present work builds upon and extends previous studies on multi-objective BO of laser wakefield accelerator configurations [24, 25], with the goal of supporting the development of the forthcoming high-repetition-rate LWFA facility at LOA. Our objective is to assess the capabilities of this new accelerator and to identify the key parameters that must be engineered to ensure its adaptability to a variety of user requirements. To this end, we employ BO to efficiently guide the exploration of PIC simulations performed with the spectral solver FBPIC [26]. This approach enables the systematic investigation of the parameter space governing electron injection and acceleration, with a particular focus on plasma density profile tuning—an essential feature for remote-controlled and automated operation. The present study extends previous optimization efforts [27, 28] by considering the full propagation of the laser and electron beam through the plasma, including the plasma exit phase, and by introducing additional parameters such as gradient injection and down-ramp length. Through this work, we aim to refine the understanding of LWFA optimization strategies and to provide practical guidelines for the design and operation of next-generation, high-repetition-rate laser wakefield accelerators.

The paper is organized as follows. In section 2, we present the physical and numerical modeling of the system, introducing the laser and plasma density parameters as well as the optimization procedures employed. In section 3, we examine the mechanisms of electron injection and acceleration under various plasma density profile configurations. Section 4 discusses the evolution of the beam emittance and divergence during plasma exit. We then analyze in section 5 the results of a start-to-end optimization. Finally, we conclude by summarizing the key findings and identifying the system components that should be prioritized in the engineering of the forthcoming LAPLACE acceleration facility.

2. Modeling

This section describes the physical and numerical modeling framework used throughout this study. We detail the laser and plasma parameters, the setup of the PIC simulations performed with the quasi-3D spectral code FBPIC [26], and the quasi-static envelope code Wake-T [29] and the implementation of the BO workflow, including its integration with GPU computing resources.

Laser and plasma configuration. The simulations were designed to reproduce the expected operating conditions of the forthcoming LAPLACE facility. The system corresponds to a Ti:Sapphire laser delivering pulses of 20 fs duration and 200 mJ energy, focused to a 10 μm FWHM spot within a supersonic gas jet, thereby generating plasma wakefields that drive electron acceleration. The plasma density profile (figure 1) consists of an ascending ramp, a density peak, a sharp injection gradient connecting the peak to the plateau, and a final down-ramp at the plasma exit. This structure allows improved stability and reliability of the injection of the electron bunch [30–33]. The height of the density peak, n_0, was fixed at 20% relative to the plateau to remain within experimentally achievable conditions. We demonstrated such results in previous work [34] and we are confident that this can be reproduced in LAPLACE.

Implementation of PIC simulations. In the first stage of the study (section 3), we performed kinetic simulations using FBPIC to investigate the injection and acceleration of electrons. The simulation domain was restricted to the light gray region in figure 1, ending at the plateau to exclude the plasma exit. The code was operated in the Lorentz-boosted frame, which significantly reduces computation time while maintaining the accuracy of the wakefield structure [35, 36]. Depending on the chosen density parameters, the total propagation length varied between 0.324 mm and 1.148 mm. The detailed simulation and laser parameters are provided in table A1.

In the subsequent phase (section 5), the plasma exit was explicitly included by extending the propagation length to 2.5 mm (dark gray region in figure 1). To keep these start-to-end simulations computationally tractable, the propagation domain was divided into two regions: the injection and early acceleration, modeled with FBPIC, and the subsequent acceleration and exit, modeled with Wake-T. Particles exhibiting a normalized longitudinal momentum $u_z = p_z/m_e c > 20$ were identified as belonging

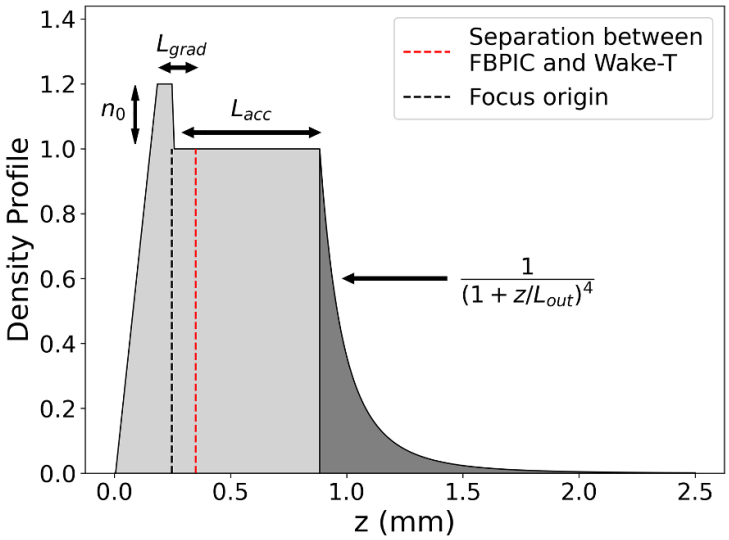

Figure 1. Plasma density profile used in simulations. The light gray region corresponds to the plasma density used in the injection study (section 3), while the dark gray extension represents the additional section included for start-to-end propagation in section 5. The red dashed line indicates the junction between the kinetic (FBPIC) and envelope (Wake-T) simulation domains.

to the injected electron bunch and subsequently transferred from one simulation code to the next. The interface between the two codes is indicated by the red dashed line in figure 1. This hybrid approach allowed us to preserve full physical fidelity during the injection phase while significantly reducing computational costs during the quasi-linear acceleration and plasma exit stages. In this part, the FBPIC structure was modeled in the laboratory reference frame. The parameters of this study are available in table A2.

Each simulation was executed on NVIDIA A100 GPU. For the first batch of simulations, using only FBPIC, individual simulations typically lasted between 20 and 30 mins. The longest batch, where 400 simulations were made over 2 GPUs working in parallel, lasted 3 d and 16 h. In the start-to-end optimization, the whole batch took 3 d and 7 h to compute the same number of simulations, having an individual mean duration of 24 mins.

BO framework. The exploration of the parameter space was guided by a BO framework implemented in Python through the Optimas library [37]. This library provides efficient parallel task scheduling across GPUs, using LibEnsemble [38] and integrates natively with the Ax API [39], which builds on the BoTorch library [40] for Gaussian-process-based optimization. The underlying principle of BO is to construct a probabilistic surrogate model of the objective function based on the observations already collected [22, 23]. This surrogate model, typically a Gaussian process [41–43], encodes correlations between sampled points and thus provides, for any unsampled input, a predictive distribution characterized by an expected value and an associated uncertainty. An acquisition function is then employed to determine the next most informative point to evaluate, balancing exploitation (sampling near regions with high predicted performance) and exploration (sampling in regions of high uncertainty to avoid missing potentially optimal areas).

In the case of multiple competing objectives, the optimization process aims to maximize the expected hypervolume defined by the set of objective values, thereby progressively extending the Pareto front that represents the trade-off boundary between the different objectives [44].

The used kernel (correlation function) was a Matérn 5/2, a standard choice in Gaussian-process modeling and a constraint of the Optimas library through the AxSingleFidelityGenerator. The acquisition function used in this work is the q–noisy expected hypervolume improvement criterion [45, 46], which is particularly well-suited to multi-objective optimization in high-dimensional parameter spaces and performs equally well for noisy as well as non-noisy data. The model hyperparameters were selected to maximize the marginal likelihood of the surrogate Gaussian process. The BO was initialized using quasi-random Sobol sequences. The initial design included 16, 24, or 32 points, depending on the dimensionality of the input parameter space.

In section 3, the BO algorithm optimized four main input parameters—electron density n_e, plateau length L_{acc}, gradient length L_{grad}, and laser focus position z_f. The objectives were the beam mean energy

E_0 and its standard deviation σ_E, which represents the energy spread. Only electrons with energies above 10 MeV and with a charge greater than 10% of the spectral maximum were included in the analysis to ensure a physically relevant estimation of beam quality. Multiple optimization runs were conducted, involving between two to four parameters, to evaluate sensitivity and trade-offs between beam energy and spread.

In this study, we deliberately focused on the generation of high-energy beams with low energy spread. Nevertheless, alternative optimization strategies could have targeted high charge and large divergence, as would be relevant, for instance, in medical irradiation contexts.

In section 5, the optimization was extended to the entire plasma, replacing the injection gradient L_{grad} with the exit length L_{out} as a free parameter. A third objective—the y-axis divergence σ'_y– was added to account for transverse beam quality after the extraction. The total propagation length was fixed to 2.5 mm, and the start-to-end simulations were carried out by combining FBPIC and Wake-T as described above.

3. Injection study

This section presents the results of the PIC simulations performed with FBPIC, optimized using BO with up to four input parameters. The simulation parameters are listed in table A1. In this stage, the plasma density profile was restricted to the light gray region of figure 1, corresponding to the injection and acceleration zones only. Figure 2(a) summarizes the outcomes of several BO runs, each consisting of one hundred simulations per free parameter. Each point in the figure thus represents one complete PIC simulation.

The initial optimization investigated two parameters: the plasma electron density n_e and the plateau length L_{acc} (green series). The focus position and the injection-gradient length were fixed at $z_f = 0$ and $L_{grad} = 50\,\mu$m, respectively. As illustrated in figure 2(b), a strong correlation was observed between L_{acc} and the electron mean energy E_0. This behavior is consistent with the physical interpretation of the plateau as the main acceleration region: a longer plateau enables the wakefield to accelerate electrons over a greater distance, thereby increasing E_0. The plateau length was therefore identified as a key tuning parameter for the future LAPLACE facility, enabling control over a wide range of achievable beam energies.

In the second optimization (red series), the injection-gradient length L_{grad} was added as a third input parameter, while keeping the focus position fixed at $z_f = 0$. However, no significant performance improvement was observed, which can be attributed to several factors. First, L_{grad} may have a limited impact in the present energy regime. Alternatively, the optimization process may not have fully converged within the 300 sampled points, or, due to the stochastic nature of BO, it may not have identified a sufficiently steep gradient for this variable compared with the others. Figure 2(c) confirms that the beam energy spread σ_E shows no systematic dependence on L_{grad}. Nevertheless, L_{grad} can influence other beam characteristics—such as emittance and charge [47, 48]— which were not part of the present optimization objectives. Moreover, the imposed constraint fixing the density peak to $n_0 = 20\%$ above the plateau, regardless of the electron density n_e, may have further reduced the sensitivity of the optimization to variations in L_{grad}.

To extend the achievable energy range, the third input parameter was then switched from L_{grad} to the laser focus position z_f (orange series), while fixing $L_{grad} = 50\,\mu$m. Moving the laser focus deeper into the plasma significantly increased the mean energy of the accelerated bunch, as the laser maintained a strong wakefield over a longer propagation distance. As shown in figure 2(d), this adjustment allowed the mean energy limit to be exceeded. However, this improvement came at the cost of beam quality: the total energy spread increased due to secondary injection events occurring later along the plateau, producing a multi-peaked energy distribution. This effect is illustrated in figure 2(e) (purple triangle on the orange Pareto front). The vertical-colored length, which represents the charge at a given energy, is related to the energy spread. Two distinct injection events are observed—one near the end of the injection peak and another close to the focal plane. The corresponding beam statistics were $E_0 = 257$ MeV and $\sigma_E = 41.9$ MeV, while isolating only the first injected population yields $E'_0 = 301$ MeV and $\sigma'_E = 29.9$ MeV. Minor discrepancies with the Pareto data arise from rerunning the case at higher temporal resolution, slightly altering the charge by 0.4 pC over 15 pC.

The four-parameter optimization (blue series), including n_e, L_{acc}, L_{grad}, and z_f, provided the most comprehensive exploration of the design space and yielded the best overall performance. The resulting Pareto front delineates the trade-off between mean energy and energy spread achievable within the LAPLACE configuration by adjusting only the plasma density profile and the laser focus. The obtained range fully covers the intended energy window (50–150 MeV) with limited energy dispersion. However, as shown in figure 2(f), the lowest-spread points in the lower-right region of the Pareto were obtained

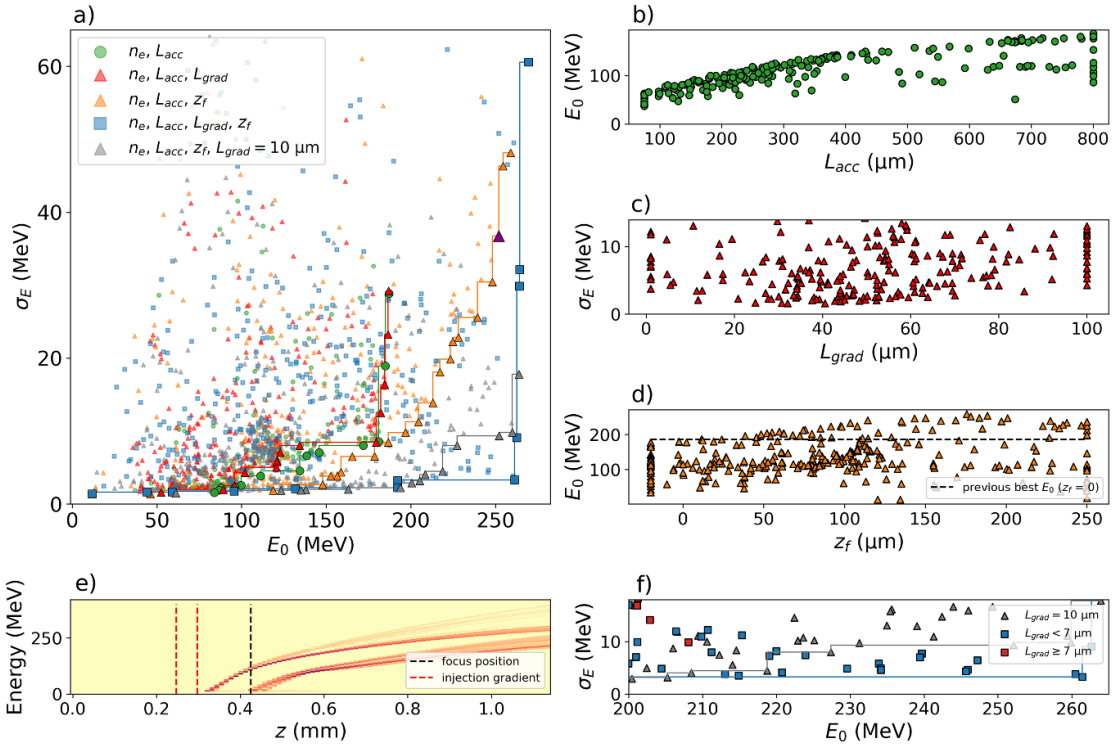

Figure 2. Bayesian optimization runs. (a) Pareto fronts for different BO runs, each using a distinct set of variable parameters. (b) Correlation between electron mean energy and plateau length. (c) Lack of correlation between energy spread and injection-gradient length. (d) Increase of the mean energy according to the position of the laser focus. (e) Example showing energy spread degradation due to multiple injections (purple triangle on the orange Pareto front). (f) Low injection-gradient cases from the four-parameter run, zoomed between the gray and blue Pareto fronts.

with extremely short injection gradients ($L_{\mathrm{grad}} < 7\,\mu$m), approaching the limit of experimental feasibility. To verify this observation, a final three-parameter optimization was conducted with L_{grad} fixed to $10\,\mu$m (gray series), using n_{e}, L_{acc}, and z_{f} as the active variables. The resulting Pareto front confirms that this configuration already spans the full energy range of interest for LAPLACE, with negligible performance loss. This finding indicates that, while L_{grad} contributes to fine-tuning the energy spread at the boundaries of the parameter space, its influence remains secondary compared with the dominant parameters n_{e}, L_{acc}, and z_{f}.

In conclusion, the injection-gradient length exerts a minor yet measurable influence on the improvement of energy dispersion, most notably in the high-energy regime, as evidenced by the comparison between the orange ($L_{\mathrm{grad}} = 50\,\mu$m) and gray ($L_{\mathrm{grad}} = 10\,\mu$m) cases in figure 2(a), with $10\,\mu$m representing the lower bound of practical [33, 34]. However, its influence is overshadowed by the dominant effects of electron density, plateau length, and laser focus position. The results also underline the balance that must be maintained between the number of free parameters and the sampling density of the optimization: increasing dimensionality without sufficient sampling can obscure meaningful correlations as in the red series. To keep the parameter space tractable and consistent with the LAPLACE target energy range, the injection-gradient length was therefore excluded from subsequent optimizations.

It is important to note that further reductions of the injection gradient length would not lead to unlimited improvements in beam quality, as other physical mechanisms ultimately impose fundamental limits. In particular, dephasing plays a central role: as the electron bunch accelerates, its longitudinal phase-space evolution leads to minimum energy spread only near the dephasing length, while deviations from this condition result in increased momentum spread. In addition, beam loading can become a limiting factor when significant charge is injected. The self-fields of the electron bunch then modify the wakefield, causing electrons at different longitudinal positions to experience different accelerating gradients, which intrinsically increases the energy spread. Signatures of beam loading were observed in several simulations of this study.

Overall, the BO framework not only enhanced the achievable beam performance but also uncovered nontrivial parameter correlations that would be difficult to identify through manual scans, demonstrating the effectiveness of ML–driven approaches for laser–plasma accelerator design.

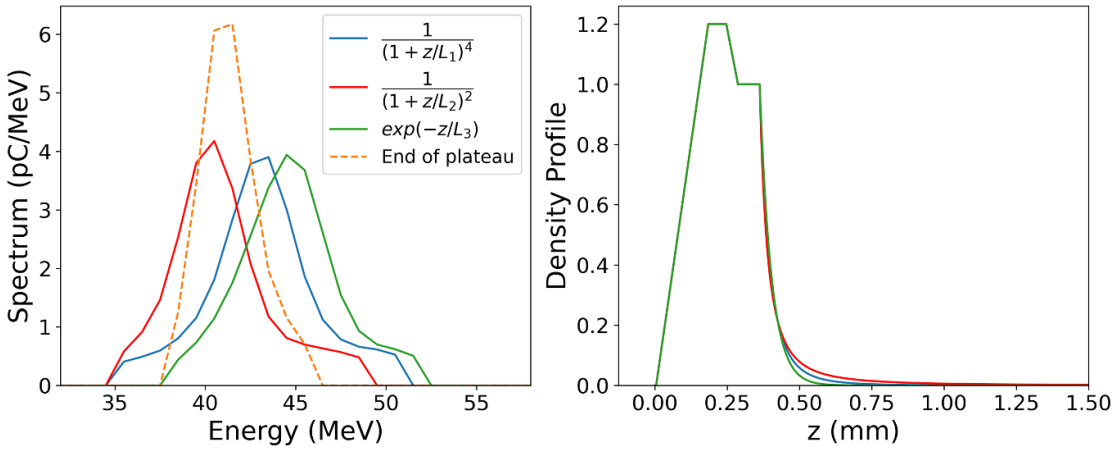

Figure 3. Comparison of energy spectra before and after the plasma exit. Right: the three down-ramp profiles used, L_2 and L_3 were chosen in order to fit the blue curve over the first 400 μm. Left: the spectrum at the end of the plateau (dashed orange) and the spectra obtained after an additional 4 mm of propagation across each respective ramp (solid curves).

4. Beam extraction stage

While section 3 demonstrated the dominant influence of the parameters n_e, L_{acc}, and z_f on electron injection and acceleration, the previous optimizations did not account for the plasma exit region, which can significantly affect the final beam quality. Figure 3 shows three fitted down-ramp profiles, each adjusted over the first 400 μm of the density transition. The corresponding characteristic lengths, obtained by optimizing the fit within this interval, are $L_1 = 133\,\mu$m, $L_2 = 53.8\,\mu$m and $L_3 = 40.4\,\mu$m. The orange dashed curve shows the spectrum at the end of the acceleration plateau; the solid curves show the spectra after an additional 4 mm of propagation through the plasma-to-vacuum transition. This example indicates that different ramp shapes and lengths may modify the electron spectrum. It is therefore important to explore the exit region since it can have a non-negligible impact on downstream beam transport and application suitability.

One of the key roles of a well-designed down-ramp is to act as a plasma lens: as the electron density decreases the focusing strength of the plasma wakefield is progressively reduced, which can lead to a controlled increase of the beam transverse size and a concomitant reduction of its divergence [49, 50]. This action can ease downstream capture by conventional optics or transport lines. The transverse quality of the bunch is typically quantified by the RMS divergence, σ_x', and the normalized emittance, $\varepsilon_{n,x}$ [51]:

$$\sigma_x' = \sqrt{\mathrm{Var}\,(x')}, \qquad \varepsilon_{n,x} = \frac{1}{m_e c}\sqrt{\mathrm{Var}\,(x)\,\mathrm{Var}\,(p_x) - \mathrm{Cov}\,(x,p_x)^2}, \tag{1}$$

where x is the transverse coordinate, p_x the transverse momentum, and $x' = \arctan(p_x/p_z)$ the particle angle. To evaluate these quantities we compute the second-order statistical moments of the particle distribution. For a set of N particles:

$$\mathrm{Var}\,(x) = \frac{1}{N}\sum_{i=1}^{N}(x_i - \mu_x)^2, \qquad \mathrm{Cov}\,(x,p_x) = \frac{1}{N}\sum_{i=1}^{N}(x_i - \mu_x)\left(p_{x,i} - \mu_{p_x}\right), \tag{2}$$

with $\mu_x = N^{-1}\sum_i x_i$ and $\mu_{p_x} = N^{-1}\sum_i p_{x,i}$, the average position and momentum. Equivalent expressions hold in the y-plane.

It is important to emphasize that even if the RMS geometric emittance is conserved in ideal linear focusing channels, the normalized emittance observed in a drift can increase due to correlations between transverse motion and energy (chromatic effects) [52, 53]. Under the standard approximations—ultra-relativistic motion ($\beta \approx 1$), small angles ($p_x \ll p_z$), weak correlations between transverse coordinates, both in positions and angles, and energy, and near-constant mean Lorentz factor γ_0– one may express the normalized emittance as:

$$\varepsilon_{n,x} \approx \gamma_0\sqrt{\sigma_e^2\,\sigma_x^2\sigma_x'^2 + \varepsilon_x^2}, \tag{3}$$

where $\sigma_x = \sqrt{\mathrm{Var}(x)}$ is the RMS beam radius, σ_x' the RMS divergence described above, $\sigma_e = \sqrt{\mathrm{Var}(\beta\gamma)}/\gamma_0$ the relative energy spread, and $\varepsilon_x = \sqrt{\mathrm{Var}(x)\mathrm{Var}(x') - \mathrm{Cov}(x,x')^2}$ the RMS geometric emittance. In a long drift where $\sigma_x(z) \simeq \sigma_x' z$, this yields:

$$\varepsilon_{n,x}(z) = \gamma_0 \sqrt{\sigma_e^2 \sigma_x'^4 z^2 + \varepsilon_x^2}, \tag{4}$$

which explicitly shows how finite energy spread σ_e and divergence σ_x' drive emittance growth during a drift. Consequently, reducing the divergence at the plasma exit mitigates subsequent emittance growth and improves transportability. We then added this parameter as an objective in table A2.

The underlying physics of this phenomenon can be described as follows: when the laser is focused into the plasma, it excites a cavity whose associated charge distribution generates a linear focusing field acting on the accelerated electron bunch. Assuming an ideal plasma channel, the resulting transverse equation of motion can be solved analytically. To investigate the beam dynamics in the density down-ramp, we employ this linear model while neglecting longitudinal fields (i.e. assuming a constant Lorentz factor γ_0), space-charge effects, and nonlinear transverse forces, and further assuming paraxial propagation. Within these conditions, the transverse displacement $x(z)$ of an electron satisfies:

$$\frac{\mathrm{d}^2 x}{\mathrm{d}z^2} + k_\beta^2(z)\, x = 0, \tag{5}$$

with the betatron wavenumber $k_\beta(z) = k_p(z)/\sqrt{2\gamma_0}$ and the plasma wavenumber $k_p(z) = \sqrt{n_e(z) e^2/(m_e \varepsilon_0)}/c$. In this regime, particles perform betatron oscillations with local period $\lambda_\beta(z) = 2\pi/k_\beta(z)$. If the ramp length allows several betatron oscillations the transition is called 'adiabatic', and the beam radius can grow while its divergence decreases without significant emittance dilution; if the ramp is too short (non-adiabatic), mismatches and phase mixing may produce emittance growth and spectral distortions [52].

In [54], Floettmann proposed a profile of the form:

$$n_e(z) = \frac{n_e}{(1 + z/L_{\mathrm{out}})^4}, \tag{6}$$

with the adiabatic criterion written as:

$$2\pi L_{\mathrm{out}} \gg \lambda_{\beta_0},$$

where λ_{β_0} is the betatron wavelength evaluated at the plateau entrance of the ramp. However, other functional forms (e.g. exponential or polynomial profiles) may also be relevant depending on the gas-flow and nozzle geometry and therefore merit investigation:

$$n_e(z) = \frac{n_e}{(1 + z/L_{\mathrm{out}})^2}, \tag{7}$$

$$n_e(z) = n_e \exp(-z/L_{\mathrm{out}}). \tag{8}$$

To compare practical options we performed particle tracking using equation (5) for three ramp shapes (equations (6)–(8)) with an ensemble of 10^4 particles initialized with $\sigma_x = 0.2\,\mu$m, $\sigma_x' = 5$ mrad, $\varepsilon_{n,x} = 294\,\mu\mathrm{m}\cdot\mathrm{mrad}$, a mean energy $E_0 = 150$ MeV, and a nominal plateau density $n_e = 10^{19}\,\mathrm{cm}^{-3}$ (giving $\lambda_{\beta_0} \approx 250\,\mu$m). The results (figure 4) indicate that equation (6) yields slightly lower final divergence for comparable characteristic lengths, but it seems only marginal as the final divergence reached respectively 4% (equation (6)), 6.6% (equation (7)) and 7% (equation (8)) of the divergence initial value. Moreover, the left side of figure 4 presents 4 domains where the optimal is alternatively in I the equation (7), in II equation (6), in III equation (8) and in IV equation (6). The compromise lies in the balance between a strong focusing power, quickly reaching low divergences like the red dashed curve and a long propagation, adiabatically reducing the forces applied on the beam, like Floettmann's suggestion.

On the basis of these considerations and of practical constraints on gas-target geometry and available propagation length, we selected Floettmann's profile (equation (6)) for the start-to-end optimization. This choice represents a compromise that provides efficient reduction of divergence without demanding excessively long ramps, while remaining achievable with realistic nozzle and vacuum-chamber designs.

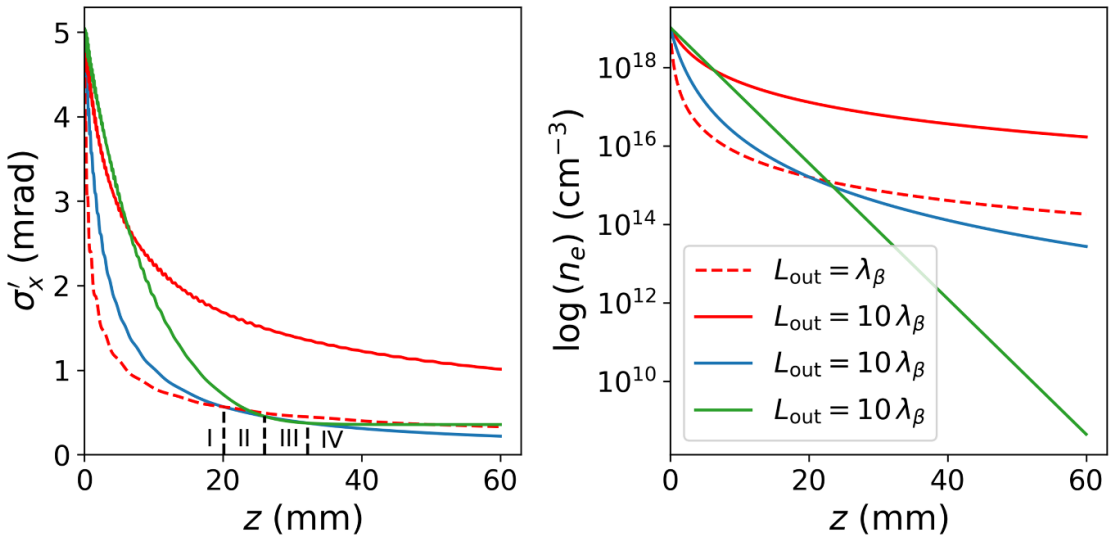

Figure 4. Divergence evolution for three different down-ramp shapes. Left: divergence evolution across propagation, the blue curve represents equation (6), the red one equation (7) and the green one equation (8). Right: plasma density profile in log-scale. Solid lines: large L_{out} (adiabatic case); dashed red: faster decay with small L_{out}.

5. Start-to-end optimization

To enable a full start-to-end optimization, we introduced the dark gray region shown in figure 1. As described in section 2, we optimized the following four parameters: electron density n_e, plateau length L_{acc}, laser focal position z_f, and exit down-ramp length L_{out}. The injection gradient was fixed to $L_{grad} = 10\,\mu$m. The three objectives were the electron mean energy and spread, E_0 and σ_E, as well as its y-axis divergence σ'_y. To capture the influence of the down-ramp, the total simulation length was doubled. To maintain computational feasibility under these extended conditions, the modeling was split between FBPIC and Wake-T. The parameter ranges are summarized in table A2, while figure 5 highlights the good agreement obtained between both codes.

The connection between the two codes is successfully achieved, and the results remain very close to those obtained using full FBPIC simulations, with only a slight overestimation of the divergence reduction. The case shown in figure 5 is one of the best optimization outcomes, yielding a mean energy of 233 MeV, an energy spread of 4 MeV, a divergence of 0.8 mrad, and a charge of 21 pC. It corresponds to simulation 128 in table 1. However, these favorable results are significantly affected in full FBPIC runs by multiple-injection effects (see section 3 and figure 2(e)), which are not captured in Wake-T simulations.

The three projections of the Pareto front are displayed in figure 6. The figure 6(c) E_0–σ_E projection exhibits an improvement over the results shown in figure 2(a), demonstrating lower energy spreads at higher energies. This improvement originates from the change in the optimization strategy. Moreover, the figure 6(a) E_0–σ'_y projection reveals an expected trend: higher-energy beams exhibit reduced divergence, since their momentum is more strongly aligned along z, decreasing $y' = \arctan(p_y/p_z)$. This results in the Pareto front being bounded on the right-hand side of the plot.

Despite the careful preparation of this start-to-end study, including the choice of exit density profile and beam quality metric at the plasma exit, we do not observe a significant influence of the parameter L_{out} on the output beam properties. Its correlation with other objectives across the full dataset is close to zero, and similar observations could be made from a plot analogous to figure 2(c). Several representative results are reported in table 1. When comparing cases with the longest and shortest exit lengths (simulations 112 and 128 versus 153 and 177), the beam divergences remain of the same order of magnitude. This is the case, for example, in simulations 128 and 153 that use nearly identical densities and focal positions, differing slightly in L_{acc} and strongly in L_{out}. Although this difference leads to different final energies, the discrepancy in divergences is of the same order of magnitude as that observed between full FBPIC simulations and hybrid FBPIC/Wake-T simulations, and thus remains within the uncertainty bounds of our study. Even though sub-mrad divergences are rare, cases like simulation 234 demonstrate that such values can be achieved with exit lengths substantially shorter than in simulation 128. The spectra of simulations 128 and 153 are shown in figure 6(b).

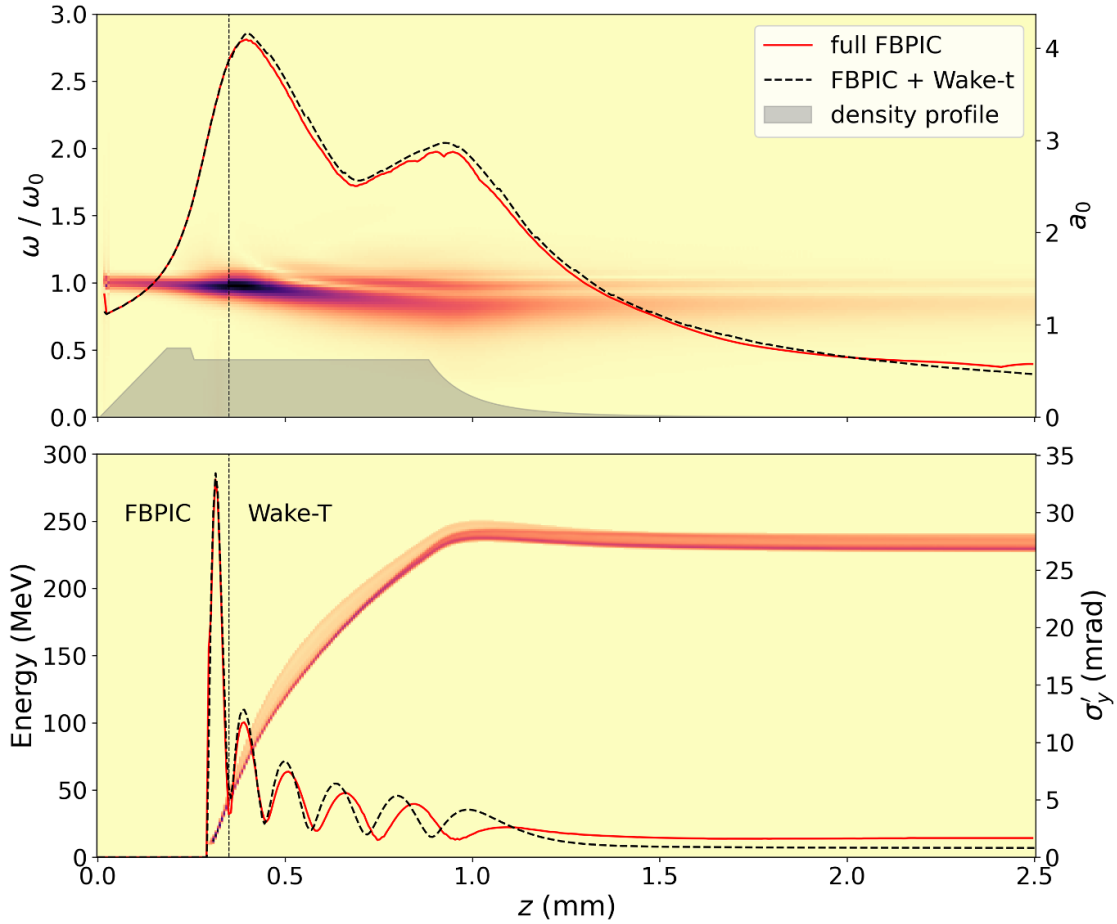

Figure 5. Evolution of the laser and electron spectra (top and bottom panels, respectively) during propagation. The left axis corresponds to the color scale: the upper panel shows the spectral dispersion of the laser fundamental frequency, and the lower panel displays the electron energy spectrum. The additional curves represent the divergence of the filtered electron beam, σ_y', and the normalized laser amplitude, a_0, whose corresponding scale is shown on the right axis. The results correspond to simulation 128 in table 1.

Table 1. Example of results obtained during the start-to-end optimization. The laser was linearly polarized along the x axis.

	Inputs				Objectives			Beam parameters					
	n_e $(\times 10^{18}\,\mathrm{cm}^{-3})$	L_{acc} (μm)	z_f (μm)	L_{out} (μm)	E_0 (MeV)	σ_E (MeV)	σ_y' (mrad)	Q (pC)	σ_x (μm)	σ_y (μm)	σ_x' (mrad)	$\varepsilon_{n,x}$ (μm mrad)	$\varepsilon_{n,y}$ (μm mrad)
sim 112	9.06	756	155	400	249	5.72	1.67	21.2	0.502	2.01	0.388	26.7	274
sim 128	9.27	626	158	400	233	3.81	0.802	21.1	0.381	1.49	0.270	24.1	244
sim 153	9.20	800	160	20	263	3.88	1.75	20.0	0.903	2.75	0.634	23.2	252
sim 177	9.80	800	155	20	255	3.87	2.79	25.4	1.45	3.80	1.04	28.0	280
sim 234	6.65	376	126	275	146	3.06	0.883	16.2	0.404	1.12	0.321	12.9	156

A possible explanation for the weak sensitivity to L_{out} is that, similar to L_{grad}, the optimization algorithm may not have sufficiently explored this parameter space. However, a more plausible interpretation is that the laser is already too depleted to sustain a well-evacuated plasma channel capable of providing transverse focusing. Consequently, the previously described model is no longer applicable. We therefore conclude that L_{out} is not intrinsically unimportant, but in the present configuration its impact is overshadowed by the dominant roles of n_e, L_{acc}, and z_f. Further investigation will be required in future studies to fully assess its relevance.

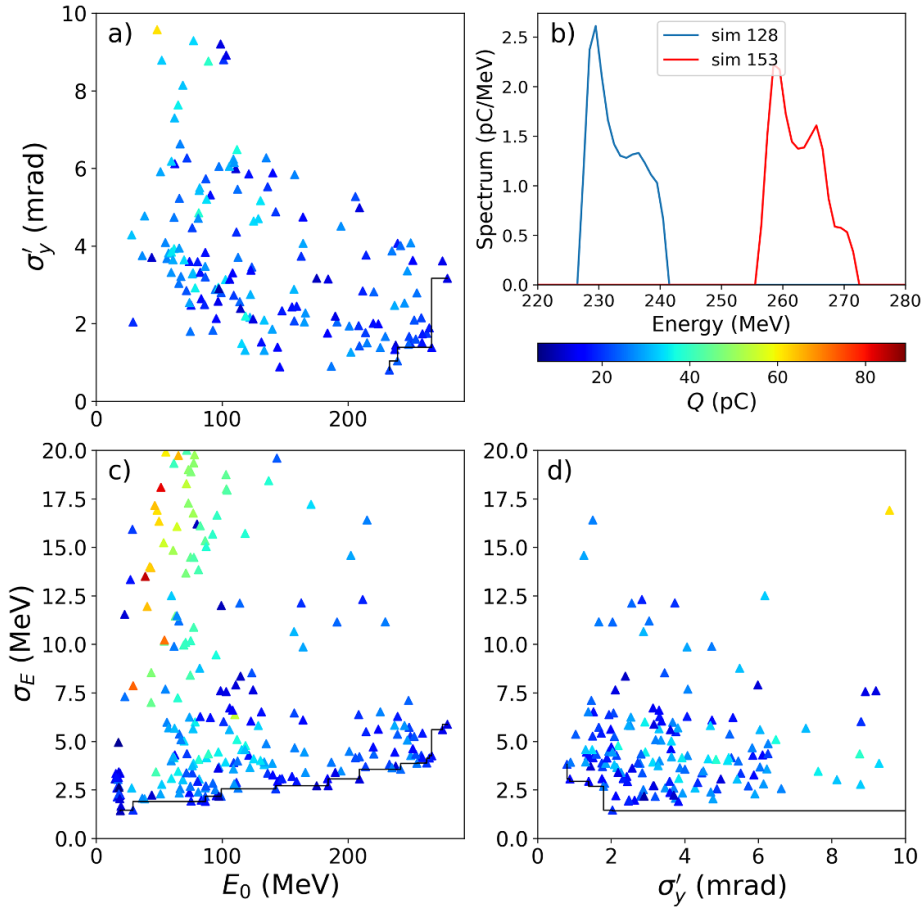

Figure 6. Projections of the Pareto surface for the start-to-end optimization, considering the mean energy E_0, energy spread σ_E, and vertical divergence σ_y' of the filtered electron beam. The color scale indicates the charge. The spectra of simulations 128 and 153 are also shown in the top right panel. (a) Projection $E_0 - \sigma_y'$ (b) example of spectra (c) Projection $E_0 - \sigma_E$ and (d) Projection $\sigma_y' - \sigma_E$. As Wake-T cannot capture secondary injection, the results should be interpreted as an upper confidence bound, but they still apply to the highest energy electron beam that does not witness the modifications of the fields behind it.

6. Conclusion

In this work, we performed a comprehensive start-to-end investigation of laser–plasma acceleration in the nonlinear blowout regime, combining high-fidelity PIC simulations with BO techniques. Our study focused on identifying the plasma-density parameters that most strongly influence electron injection, acceleration, and transport toward a conventional beamline.

We demonstrated that the final beam energy and energy spread are predominantly governed by the plasma density, the length of the acceleration plateau, and the laser focal position relative to the wake-field structure. These parameters exhibit strong nonlinear correlations, which the optimization framework successfully exploited to improve both metrics, yielding impressive results even with modest laser parameters. Introducing additional variables improved the Bayesian convergence, but this benefit must be balanced against the available number of sampling evaluations.

We further showed that the plasma-to-vacuum transition—often neglected in optimization studies—can have a significant impact on the preservation of beam emittance and spectral quality. A controlled density down-ramp allowing adiabatic relaxation of the focusing forces was found to be advantageous but not strictly essential for our parameters: depending on the profile, it can help maintain low divergence and limit emittance growth induced by phase-space decoherence. We also demonstrated that long-distance propagation can be included within the optimization loop, achieving an optimization over 2.5 mm with 400 simulations in a reasonable computational time.

Among the additional parameters explored, the injection-gradient length appeared largely dominated by the main drivers over the central region of the design space, but offered beneficial tuning capability near the edges of the parameter domain. It is important to stress that this result was obtained by fixing the height of the density peak relative to the plateau, although these parameters are strongly correlated. In contrast, the exit down-ramp length showed no clear correlation with any of the targeted objectives within the present run, which prevents drawing a definitive conclusion about its optimal value.

Future studies could compare Bayesian-optimization convergence for different classes of down-ramp profiles—possibly reducing the dimensionality of the search—to better assess the influence of the exit-transition length. The interplay between the injection gradient and exit-region metrics such as emittance, divergence, or beam radius also deserves further investigation. Finally, as LAPLACE will be extended to an energy level of 1 J, future study may also be interested in higher laser energy behavior.

Overall, this work establishes a robust methodology for the design and optimization of next-generation plasma accelerators and identifies key elements to prioritize in the realization of a reliable user-oriented facility.

Acknowledgments

This project has received funding from CNRS, Région Ile-de-France and Ministère de l'Enseignement supérieur et de la Recherche through the CPER LAPLACE project. We acknowledge the support of the Institut Pierre Lamoure for the chaire *Accélération laser-plasma haute cadence*.

Data availability statement

The data that support the findings of this study are available from the contact author upon reasonable request. The data cannot be made publicly available upon publication because no suitable repository exists for hosting data in this field of study.

Appendix. Parameter tables

This appendix compiles the characteristic parameter tables associated with the optimization studies presented in this work. It provides a detailed summary of the input parameters, laser configuration, numerical simulation settings, and objective quantities considered in the analysis.

Table A1. Optimization parameters of section 3.

Input parameters				
Parameter	symbol	value	unit	
Electron density	n_e	[4 ; 15]	$\times 10^{18}\,\mathrm{cm}^{-3}$	
Plateau length	L_{acc}	[75 ; 800]	μm	
Injection-gradient length	L_{grad}	[1 ; 100]	μm	
Laser focus position[a]	z_f	[−20 ; 250]	μm	
Laser parameters				
Parameter	symbol	value	unit	
Wavelength	λ_0	800	nm	
Energy	E_{laser}	200	mJ	
Duration (FWHM)	τ_0	20	fs	
Diameter (FWHM)	D_0	10	μm	
Waist	w_0	8.49	μm	
Rayleigh length	z_R	283	μm	
Vector potential	a_0	1.97		
Intensity	I_0	8.29	$\times 10^{18}\,\mathrm{Wcm}^{-2}$	
Critical density	n_c	1.74	$\times 10^{21}\,\mathrm{cm}^{-3}$	
Laser polarization	θ_{pol}	0	rad	(linear along $\vec{e}_x$)
FBPIC parameters				
Parameter	symbol	value	unit	formula
Plasma start	p_{zmin}	0	mm	
Plasma end	p_{zmax}	[0.324 ; 1.148]	mm	$0.2475 + L_{acc} + L_{grad}$
Number of modes	N_m	3		
Particle per cell	$p_{n,z}$; $p_{n,r}$; $p_{n,t}$	3 ; 3 ; 16		
Pic injection	n_0	0.2		
Boost factor	γ_{boost}	2.5		
Laser centroid position	z_0	−17.8	μm	$p_{zmin} - 3.5\,c\tau_0$
Laser maximal radius	R	[11.2 ; 17.6]	μm	$w_0\left(1+(z_0-z_f)^2/z_R^2\right)^{1/2}$
Plasma frequency	ω_p	[1.13 ; 2.18]	$\times 10^{14}\,\mathrm{rad\,s}^{-1}$	$(n_e e^2/m_e \varepsilon_0)^{1/2}$
Plasma wavelength	λ_p	[8.62 ; 16.7]	μm	$2\pi c/\omega_p$
z-step	dz	25	nm	$\lambda_0/32$
r-step	dr	112.5	nm	4.5 dz
Time-step	dt	min(74.9 ; 83, 4)	as	$\min\left(\frac{\mathrm{d}r}{2c\gamma_b}\,;\,\frac{\mathrm{d}z}{c}\right)$
Window width	r_{max}	[36 ; 56.4]	μm	$3.2\,R$
Window length	L_z	[34.4 ; 47.2]	μm	$4c\tau_0 + 1.6\,\lambda_p$
Grid copropagative speed	v_{co}	−2.75	$\times 10^8\,\mathrm{m\,s}^{-1}$	$-c\left(1-1/\gamma_b^2\right)^{1/2}$
Objectives				
Parameter	symbol	unit		
Mean electron energy	E_0	MeV		
Electron energy spread	σ_E	MeV		

[a] The focus origin is set at the beginning of the injection gradient (black dashed line of figure 1).

Table A2. Optimization parameters of section 5, laser parameters were identical to those in table A1 and few changes were made to the FBPIC script.

Input parameters			
Parameter	symbol	min	unit
Electron density	n_e	[4 ; 15]	$\times 10^{18}\mathrm{cm}^{-3}$
Plateau length	L_{acc}	[75 ; 800]	μm
Laser focus position[a]	z_f	[−20 ; 250]	μm
Down-ramp length	L_{out}	[20 ; 400]	μm

FBPIC changes				
Parameter	symbol	value	unit	formula
Injection-gradient length	L_{grad}	10	μm	
Plasma end	p_{zmax}	2.5	mm	
End FBPIC	$z_{end,\,PIC}$	350	μm	
Number of modes	N_m	2		
Particle per cell	$p_{n,z}$; $p_{n,r}$; $p_{n,t}$	2 ; 2 ; 8		
Boost factor	γ_{boost}	0		
r-step	dr	100	nm	4 dz
Time-step	dt	83.4	as	dz/c
Window width	r_{max}	[31.5 ; 47.7]	μm	
Window length	L_z	37.6	μm	
Grid copropagative speed	v_{co}	0		

Wake-T parameters				
Parameter	symbol	value	unit	formula
Length	$L_{Wake\text{-}T}$	2.15	mm	$p_{zmax} - z_{end,\,PIC}$
Window width	$r_{max,\,Wake\text{-}T}$	[189 ; 286.2]	μm	6 r_{max}
Window length	$L_{z,Wake\text{-}T}$	37.6	μm	L_z
Grid r	n_r	1908		4 $r_{max}/$dr
Grid ξ	n_ξ	752		$L_z/(2\,$d$z)$
Step before updating wakefield	dz_{fields}	6.35	μm	0.15 L_z
Particle per radial cell	ppc	8		
Background limiting gamma	γ_{max}	25		
Laser substeps	2			
Laser envelope grid	$n_{\xi,laser}$; $n_{r,laser}$	752 ; 954		$L_z/(2\,$d$z)$; 2 $r_{max}/$dr

Objectives		
Parameter	symbol	unit
Mean electron energy	E_0	MeV
Electron energy spread	σ_E	MeV
Y-axis divergence	σ'_y	mrad

[a] The focus origin is set at the beginning of the injection gradient (black dash line of figure 1).

ORCID iDs

Semion Tchetovsky 0009-0009-8111-1864
Igor A Andriyash 0000-0003-0313-4496
Jérôme Faure 0000-0003-0538-5774

References

[1] Tajima T and Dawson J M 1979 Laser electron accelerator *Phys. Rev. Lett.* **43** 267–70
[2] Esarey E, Schroeder C B and Leemans W P 2009 Physics of laser-driven plasma-based electron accelerators *Rev. Mod. Phys.* **81** 1229–85
[3] Malka V 2012 Laser plasma accelerators *Phys. Plasmas* **19** 055501
[4] 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
[5] Mangles S P D *et al* 2004 Monoenergetic beams of relativistic electrons from intense laser–plasma interactions *Nature* **431** 535–8

[6] Geddes C G R, Toth C, van Tilborg J, 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
[7] Leemans W P, Nagler B, Gonsalves A J, Tóth C, Nakamura K, Geddes C G R, Esarey E, Schroeder C B and Hooker S M 2006 GeV electron beams from a centimetre-scale accelerator *Nat. Phys.* **2** 696–9
[8] Wang X *et al* 2013 Quasi-monoenergetic laser-plasma acceleration of electrons to 2 GeV *Nat. Commun.* **4** 1988
[9] Kim H T *et al* 2017 Stable multi-GeV electron accelerator driven by waveform-controlled PW laser pulses *Sci. Rep.* **7** 10203
[10] Gonsalves A *et al* 2019 Petawatt laser guiding and electron beam acceleration to 8 GeV in a laser-heated capillary discharge waveguide *Phys. Rev. Lett.* **122** 084801
[11] Oubrerie K *et al* 2022 Controlled acceleration of GeV electron beams in an all-optical plasma waveguide *Light Sci. Appl.* **11** 180
[12] Miao B *et al* 2022 Multi-GeV electron bunches from an all-optical laser wakefield accelerator *Phys. Rev. X* **12** 031038
[13] Picksley A *et al* 2024 Matched guiding and controlled injection in dark-current-free, 10-GeV-class, channel-guided laser-plasma accelerators *Phys. Rev. Lett.* **133** 255001
[14] Maier A R *et al* 2020 Decoding sources of energy variability in a laser-plasma accelerator *Phys. Rev. X* **10** 031039
[15] Rovige L *et al* 2020 Demonstration of stable long-term operation of a kilohertz laser-plasma accelerator *Phys. Rev. Accel. Beams* **23** 093401
[16] Monzac J *et al* 2024 Optical ionization effects in kHz laser wakefield acceleration with few-cycle pulses *Phys. Rev. Res.* **6** 043099
[17] Laser plasma acceleration center 2025 (available at: https://laplace-loa.fr/)
[18] Döpp A, Eberle C, Howard S, Irshad F, Lin J and Streeter M 2023 Data-driven science and machine learning methods in laser–plasma physics *High Power Laser Sci. Eng.* **11** e55
[19] Shalloo R J *et al* 2020 Automation and control of laser wakefield accelerators using Bayesian optimization *Nat. Commun.* **11** 6355
[20] Duris J *et al* 2020 Bayesian optimization of a free-electron laser *Phys. Rev. Lett.* **124** 124801
[21] Jalas S, Kirchen M, Messner P, Winkler P, Hübner L, Dirkwinkel J, Schnepp M, Lehe R and Maier A R 2021 Bayesian optimization of a laser-plasma accelerator *Phys. Rev. Lett.* **126** 104801
[22] Jones D R, Schonlau M and Welch W J 1998 Efficient global optimization of expensive black-box functions *J. Glob. Optim.* **13** 455–92
[23] Frazier P I 2018 A tutorial on Bayesian optimization
[24] Irshad F, Karsch S and Döpp A 2023 Multi-objective and multi-fidelity Bayesian optimization of laser-plasma acceleration *Phys. Rev. Res.* **5** 013063
[25] Irshad F, Eberle C, Foerster F M, v. 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** 085001
[26] Lehe R, Kirchen M, Andriyash I A, Godfrey B B and Vay J-L 2016 A spectral, quasi-cylindrical and dispersion-free particle-in-cell algorithm *Comput. Phys. Commun.* **203** 66–82
[27] Valenta P, Esirkepov T, Ludwig J, Wilks S and Bulanov S 2025 Bayesian optimization of electron energy from laser wakefield accelerators *Phys. Rev. Accel. Beams* **28** 094601
[28] Gardner C and Dollar F 2025 Implementing Bayesian optimization to compare injection regimes and pulse propagation in two stage LWFA *New J. Phys.* **27** 124301
[29] Pousa A F, Assmann R and de la Ossa A M 2019 Wake-T: a fast particle tracking code for plasma-based accelerators *J. Phys.: Conf. Ser.* **1350** 012056
[30] Suk H, Lee H J and Ko I S 2004 Generation of high-energy electrons by a femtosecond terawatt laser propagating through a sharp downward density transition *J. Opt. Soc. Am. B* **21** 1391–6
[31] Brantov A V, Esirkepov T Z, Kando M, Kotaki H, Bychenkov V Y and Bulanov S V 2008 Controlled electron injection into the wake wave using plasma density inhomogeneity *Phys. Plasmas* **15** 073111
[32] Faure J, Rechatin C, Lundh O, Ammoura L and Malka V 2010 Injection and acceleration of quasimonoenergetic relativistic electron beams using density gradients at the edges of a plasma channel *Phys. Plasmas* **17** 083107
[33] 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. ST Accel. Beams* **13** 091301
[34] Rovige L *et al* 2021 Symmetric and asymmetric shocked gas jets for laser-plasma experiments *Rev. Sci. Instrum.* **92** 083302
[35] Kirchen M, Lehe R, Godfrey B B, Dornmair I, Jalas S, Peters K, Vay J-L and Maier A R 2016 Stable discrete representation of relativistically drifting plasmas *Phys. Plasmas* **23** 100704
[36] Lehe R, Kirchen M, Godfrey B B, Maier A R and Vay J-L 2016 Elimination of numerical Cherenkov instability in flowing-plasma particle-in-cell simulations by using Galilean coordinates *Phys. Rev. E* **94** 053305
[37] Pousa A F, Jalas S, Kirchen M, de la Ossa A M, Thévenet M, Hudson S, Larson J, Huebl A, Vay J-L and Lehe R 2023 Bayesian optimization of laser-plasma accelerators assisted by reduced physical models *Phys. Rev. Accel. Beams* **26** 084601
[38] 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
[39] Olson M *et al* 2025 Ax: a platform for adaptive experimentation *Automl 2025 Abcd Track*
[40] Balandat M, Karrer B, Jiang D R, Daulton S, Letham B, Wilson A G and Bakshy E 2020 BOTORCH: a framework for efficient Monte-Carlo Bayesian optimization *Proc. 34th Int. Conf. on Neural Information Processing Systems, NIPS'20* (Curran Associates Inc) pp 21524–38
[41] Williams C K I and Rasmussen C E 1995 Gaussian processes for regression *Proc. 9th Int. Conf. on Neural Information Processing Systems, NIPS' 95* (MIT Press) pp 514–20
[42] Rasmussen C E and Williams C K I 2006 *Gaussian Processes for Machine Learning* (The MIT Press)
[43] Görtler J, Kehlbeck R and Deussen O 2019 A visual exploration of Gaussian processes *Distill* **4** e17
[44] Roussel R, Hanuka A and Edelen A 2021 Multiobjective Bayesian optimization for online accelerator tuning *Phys. Rev. Accel. Beams* **24** 062801
[45] Daulton S, Balandat M and Bakshy E 2020 Differentiable expected hypervolume improvement for parallel multi-objective Bayesian optimization *Proc. 34th Int. Conf. on Neural Information Processing Systems, NIPS'20* (Curran Associates Inc) pp 9851–64
[46] Daulton S, Balandat M and Bakshy E 2021 Parallel Bayesian optimization of multiple noisy objectives with expected hypervolume improvement *Proc. 35th Int. Conf. on Neural Information Processing Systems, NIPS'21* (Curran Associates Inc) pp 2187–200
[47] 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

[48] Ekerfelt H, Hansson M, González I G, Davoine X and Lundh O 2017 A tunable electron beam source using trapping of electrons in a density down-ramp in laser wakefield acceleration *Sci. Rep.* **7** 12229
[49] Doss C *et al* 2019 Laser-ionized, beam-driven, underdense, passive thin plasma lens *Phys. Rev. Accel. Beams* **22** 111001
[50] Chang Y-Y, Cabadağ J C, Debus A, Ghaith A, LaBerge M, Pausch R, Schöbel S, Ufer P, Schramm U and Irman A 2023 Reduction of the electron-beam divergence of laser wakefield accelerators by integrated plasma lenses *Phys. Rev. Appl.* **20** L061001
[51] Floettmann K 2003 Some basic features of the beam emittance *Phys. Rev. ST Accel. Beams* **6** 034202
[52] Mehrling T, Grebenyuk J, Tsung F S, Floettmann K and Osterhoff J 2012 Transverse emittance growth in staged laser-wakefield acceleration *Phys. Rev. ST Accel. Beams* **15** 111303
[53] Migliorati M, Bacci A, Benedetti C, Chiadroni E, Ferrario M, Mostacci A, Palumbo L, Rossi A R, Serafini L and Antici P 2013 Intrinsic normalized emittance growth in laser-driven electron accelerators *Phys. Rev. ST Accel. Beams* **16** 011302
[54] Floettmann K 2014 Adiabatic matching section for plasma accelerated beams *Phys. Rev. ST Accel. Beams* **17** 054402