

High-brightness, symmetric electron bunch generation in a plasma wakefield accelerator via a radially-polarized plasma photocathode

J. Chappell^{1,†}, E. Archer^{1,‡}, R. Walczak^{1,2} and S. M. Hooker^{1,*}

¹*John Adams Institute for Accelerator Science and Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford, United Kingdom*

²*Somerville College, Woodstock Road, Oxford, United Kingdom*



(Received 19 May 2025; accepted 11 September 2025; published 3 October 2025)

The plasma photocathode has previously been proposed as a source of ultrahigh-brightness electron bunches within plasma accelerators. Here, the scheme is extended by using a radially-polarized ionizing laser pulse to generate high-charge, high-brightness electron bunches with *symmetric* transverse emittance. Efficient start-to-end modeling of the scheme, from ionization and trapping until drive bunch depletion, enables a multiobjective Bayesian optimization routine to be performed to understand the performance of the radially-polarized plasma photocathode, quantify the stability of the scheme, and explore the fundamental relation between the witness bunch charge and its emittance. Comparison of plasma photocathodes driven by radially- and linearly-polarized laser pulses shows that the former yields higher-brightness electron bunches when operating in the optimally-loaded regime.

DOI: 10.1103/hqqg-7rcd

I. INTRODUCTION

Light sources such as x-ray free electron lasers (XFELs) [1,2] and synchrotrons [3,4] have transformed our understanding as essential tools for studies of biological, chemical, and high energy density science. Generating this radiation requires the highest quality electron bunches, achieved by maximizing their charge density in transverse and longitudinal phase space. The metric used to quantify bunch quality is the 6D *brightness*, calculated via the equation [5]:

$$B_{6D} = \frac{I_{pk}}{\epsilon_{n,x}\epsilon_{n,y}\epsilon_{n,z}}, \quad (1)$$

where I_{pk} is the peak current of the bunch, $\epsilon_{n,x/y}$ is the transverse normalized emittance in the x/y plane, and $\epsilon_{n,z}$ is the normalized longitudinal emittance of the bunch. State-of-the-art XFEL facilities, powered by radio-frequency (rf) accelerating cavities, can achieve projected—i.e., calculated over the entire bunch—6D brightnesses of $B_{6D} \sim \mathcal{O}(1 \times 10^{24} \text{ A m}^{-3})$, with peak longitudinal slice

brightness exceeding $B_{6D} \sim \mathcal{O}(1 \times 10^{25} \text{ A m}^{-3})$ [6]. The use of rf accelerating cavities with field gradients typically in the range 10–100 MV m^{−1} at such facilities requires extended linear accelerator lengths to achieve the multi-GeV electron energies needed to generate radiation with the $\lambda_r \sim \text{nm}$ wavelengths of interest.

Plasma accelerators [7] offer a compelling alternative route to future-generation light sources thanks to the extreme accelerating gradients they can support (1–100 GV m^{−1}), enabling construction of high-quality, tunable light sources at much reduced cost and within a significantly more compact footprint compared to conventional accelerator facilities, potentially enabling wider access to such sources. Furthermore, beam-driven plasma accelerators [8,9] offer the potential to transform already existing light sources, either by acting as a “plasma booster” to increase the energy reach of such facilities [10,11], or through the generation of electron bunches with unique parameters, currently inaccessible to conventional rf accelerator technology.

One such example of the latter case is the plasma photocathode [12], a scheme in which an intense, ultrashort laser pulse co-propagates within the wakefield driven by a high-density electron bunch. Around its focus, the laser pulse selectively ionizes a high intensity threshold (HIT) dopant gas and releases electrons into the wake, which can then be trapped and accelerated. A key advantage of this scheme compared to traditional photoinjectors in rf accelerators is that the electrons are released directly into the extreme longitudinal fields provided by the plasma wakefield, providing rapid acceleration and minimizing space-charge effects that cause significant emittance growth.

*Contact author: simon.hooker@physics.ox.ac.uk

†Present address: Met Office, Exeter, Devon, United Kingdom.

‡Present address: Deutsches Elektronen-Synchrotron DESY, Notkestraße 85, 22607 Hamburg, Germany.

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International license*. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

Additionally, tuning the laser pulse intensity to be just above the ionization threshold of the HIT dopant releases electrons within a small volume and can minimize their thermal emittance [13]. Previous simulation studies [5,14,15] of the plasma photocathode scheme have indicated 6D brightness values in excess of $1 \times 10^{27} \text{ A m}^{-3}$, a significant performance increase over the bunches provided by state-of-the-art rf facilities, in only a fraction of the footprint.

However, the generation of electron beams with such exceptional quality via the plasma photocathode scheme requires low bunch charge ($\lesssim 1 \text{ pC}$) to minimize the transverse emittance and generate such extreme brightness. This leads to significant growth of the energy spread of the electron bunch during acceleration over extended distances due to the variation of the longitudinal field amplitude over the length of the bunch itself. This makes transport of the bunch outside of the plasma accelerator particularly challenging due to the inherent chromaticity of traditional transport optics, leading to emittance growth and decreased bunch brightness. An extension of the original plasma photocathode scheme [15] proposed injection of a second, higher charge and lower quality bunch to “escort” the high-quality witness bunch. The current of the escort bunch located at the rear of the wakefield would perturb the trajectories of inflowing plasma electrons, modifying its structure via beam loading [16]. By appropriately tailoring the charge distribution of the escort bunch, it was shown that it is possible to minimize the variation of the longitudinal wakefield amplitude in the region of the high-quality witness bunch, maintaining its energy spread throughout acceleration. However, this scheme involves significantly increased experimental complexity.

An alternative approach would be to increase the current of the original witness bunch such that it is capable of optimally loading the longitudinal wakefield itself. This would require increasing the amount of charge released into the wakefield by the ionizing laser pulse by increasing its intensity, which can have a detrimental effect on the quality of the electron bunch. Owing to conservation of transverse canonical momentum, electrons released via field ionization gain a transverse momentum equal to the laser vector potential at the moment of ionization—momentum that is directed in the plane of polarization. A linearly-polarized plasma photocathode therefore creates electron bunches with larger emittance in the polarization plane—an effect that grows with increasing intensity. Such an emittance asymmetry can lead to emittance mixing over extended propagation distances [17], an effect particularly prevalent for high-charge, low-emittance electron bunches, the extreme space-charge fields of which can promote strong ion motion and generate nonlinearly coupled transverse wakefields. This can lead to degradation of the electron bunch quality, decreasing its brightness.

In this paper, we overcome this limitation through use of a radially-polarized plasma photocathode. For this polarization, the laser field in a given transverse plane is

always parallel (or antiparallel) to the radial vector, which causes electrons produced via laser ionization to have an azimuthally-symmetric emittance. In contrast, for linearly- or circularly-polarized pulses, the instantaneous laser field is unidirectional in a given transverse plane, and hence is not azimuthally symmetric. Although use of a radially-polarized plasma photocathode leads to larger emittance at the point of ionization, we show that in the high-charge regime—where the final emittance of the witness bunch is dominated by emittance growth during matching of the bunch to the transverse wakefield—the average transverse emittance is minimized, and hence 6D brightness maximized, by using a radially-polarized ionizing laser pulse. Radially-polarized pulses are readily produced experimentally using various techniques such as interferometry [18,19] and segmented wave plates [20,21] and have been produced with relativistic intensities [22,23].

Use of a radially-polarized ionizing laser pulse allows high-charge, low-energy spread operation and maximizes brightness within a single injection event. Furthermore, through use of a multiobjective Bayesian optimization routine we explore the fundamental relationship between the bunch quality and charge in the plasma photocathode scheme, study the sensitivity of the optimally-loaded regime to relative timing fluctuations, and discuss its prospects as an injector for future plasma-based accelerator facilities.

II. THE RADIALLY-POLARIZED PLASMA PHOTOCATHODE

The concept of the radially-polarized plasma photocathode matches that of a linearly-polarized scheme, except for the properties of the ionizing laser pulse. In this study, a radially-polarized vortex-based Laguerre-Gaussian laser pulse [24] (denoted RPLP) is generated in simulation using the LASY [25] library. The pulse is constructed via the superposition of right- and left-handed circularly-polarized vortex beams with Laguerre-Gaussian (LG) profiles of azimuthal order $l = \pm 1$ and radial order $n = 0$. The full-width at half-maximum duration of the ionizing laser pulse was set to 30 fs, while its central wavelength was $\lambda = 800 \text{ nm}$. For this proof-of-concept simulation the spot size of the RPLP was $w_0 = 4.5 \text{ }\mu\text{m}$, giving a Rayleigh range of $z_R = 80 \text{ }\mu\text{m}$ and the pulse contained a total energy of $E_{\text{Laser}} = 0.5 \text{ mJ}$. The injection process, which occurs within the first 1 mm of propagation of the drive bunch, is simulated using FBPIC [26–28], a quasi-3D particle-in-cell (PIC) code, using $N_m = 3$ azimuthal modes. After extraction of the resulting drive and witness bunches from this simulation, their subsequent evolution is then simulated using Wake-T [29] until the point of drive bunch depletion ($z = 90 \text{ mm}$). The RPLP is not included in the Wake-T simulation since it diffracts sufficiently by $z = 1 \text{ mm}$ ($\sim 10z_R$) that its fields no longer ionize the dopant nor modify the wakefield.

Within this study, preionized hydrogen of on-axis density $n_0 = 1 \times 10^{17} \text{ cm}^{-3}$ is assumed. This gives a wakefield cavity length, $\lambda_{p0} = 2\pi c/\omega_{p0} = 105 \text{ }\mu\text{m}$ —where $\omega_{p0} = \sqrt{(n_0 e^2)/(m_e \epsilon_0)}$ —that matches the minimum electron bunch lengths available at current beam-driven plasma accelerator facilities [30,31], and thus reduces the computational cost of the simulation. We note that the plasma photocathode scheme can be used over a range of plasma densities, and the operating density is often chosen according to the drive electron bunch parameters. Higher density operation ($n_0 \sim 1 \times 10^{18} \text{ cm}^{-3}$) is made possible through the use of ultrashort (few fs) drive bunches generated in a laser-wakefield accelerator—akin to the “hybrid” plasma accelerator scheme currently being explored [32,33]—and gives access to the highest amplitude accelerating gradients as the maximum field supported by a cold plasma scales as $E_0 \propto \sqrt{n_0}$. Lower density operation ($n_0 \sim 10^{15} - 10^{16} \text{ cm}^{-3}$) is also possible using longer drive bunches (few hundred fs) available from rf accelerator facilities and is often preferred for future linear collider designs [34] due to improved drive bunch stability, while still maintaining accelerating gradients exceeding 1 GV m^{-1} .

Here, a matched, high-density, ultrarelativistic ($\gamma_0 = 2055$) electron bunch drives a blown-out wakefield. The RPLP follows closely, $\Delta\zeta = -45 \text{ }\mu\text{m}$ where $\zeta = z - ct$ is the co-moving coordinate, and locally ionizes the 5% neutral helium dopant over a small region around its focal position ($z = z_{\text{foc}} = 0.2 \text{ mm}$), releasing a significant amount of charge—initially at rest in the lab frame—into the plasma cavity. When the focal intensity of the RPLP is sufficiently high, electrons are ionized from both the first and second levels of the helium dopant, as shown by the blue and purple regions, respectively in Fig. 1, with a reduced ionization volume of the second level due to its increased ionization threshold. The ionization energies for the first and second levels of helium are 24.6 and 54.4 eV, respectively. The electrons initially slip backward in the co-moving frame of the wakefield and are rapidly accelerated by the longitudinal wakefield, leading to trapping of a fraction at the rear of the accelerating cavity.

Use of a RPLP necessitates a far-field intensity profile of a higher-order LG mode, as illustrated in the upper inset of Fig. 1, in contrast to the fundamental Gaussian profile considered in previous studies. Unlike a fundamental Gaussian pulse whose ponderomotive force pushes ionized electrons away from the axis, the “donutlike” intensity profile of the RPLP peaks off axis generating a ponderomotive force which is radially focusing for ionized electrons released within the peak intensity contour, and defocusing for those outside [35,36]. This acts to focus these electrons toward the axis, which, in combination with the focusing wakefield, results in a strongly peaked spatial distribution of trapped electrons as seen in Fig. 2(a). This ponderomotive “focusing” feature is not unique to radial




FIG. 1. Snapshot from PIC simulation demonstrating the injection of dopant electrons into the wakefield driven by a high-density, ultrarelativistic electron bunch. This timestep corresponds to the focal location of the RPLP ($z = z_{\text{foc}}$). Upper inset: the transverse intensity profile of the plasma photocathode at focus. Lower inset: the on-axis longitudinal electric field at $z = z_{\text{foc}}$ (black) and $z = 1 \text{ mm}$ (gray, dashed), including the current profile of the trapped charge ($z = 1 \text{ mm}$). The electron number density, n_e , is colored according to ionization source: preionized hydrogen (gray), He (green/blue), and He^+ (pink/purple) and normalized to the background plasma density, $n_0 = 1 \times 10^{17} \text{ cm}^{-3}$. The normalized drive bunch density, n_d , is represented by the orange/black colormap, while the intensity profile of the RPLP is shown in red.

polarization as it is due to the intensity distribution at focus and hence would also be present for a linearly-polarized higher-order LG laser pulse.

The radial nature of the electric field at focus leads to the generation of *symmetric* transverse phase-space profiles of the injected electron bunch. This is demonstrated in Fig. 2(a), which shows the simulated transverse phase space of the witness electron bunch after cessation of the injection process ($z = 1 \text{ mm}$). The distributions of trapped electrons in the




FIG. 2. Transverse phase space of the trapped electrons at (a) $z = 1 \text{ mm}$ and (b) $z = 90 \text{ mm}$, the depletion length of the drive bunch. The trapped charge from each ionization source (He: blue, He^+ : pink) is displayed above the subfigures. Logarithmic colorbars are used to improve visibility of the strongly-peaked distributions.

$x - x'$ and $y - y'$ phase spaces are identical and remain so throughout propagation as demonstrated in Fig. 2(b) ($z = 90$ mm).

Around its focus, the tightly-focused RPLP generates a significant longitudinal field as demonstrated in the lower inset of Fig. 1 at $z = z_{\text{foc}}$. This field acts to longitudinally accelerate the ionized electrons at their moment of release, but its magnitude is significantly smaller than that of the transverse laser field ($|E_x/E_z| \sim 50$) for the focusing geometry considered in this study. The initial momenta of the ionized electrons are therefore predominantly transverse, as in the linearly-polarized case, and no significant modification of the amount of charge trapped was observed for identical laser pulses with different polarization states but the same transverse intensity profile.

Figure 3 demonstrates the witness bunch parameters after depletion of the drive bunch at $z = 90$ mm. The witness bunch has a charge of 38.5 pC and an rms duration of 3.1 fs. The ultrashort bunch duration, despite the range of initial release positions of electrons in ζ , is a key feature of the plasma photocathode scheme [5] and results in a high peak current of $I_{\text{pk}} = 5.2$ kA. The longitudinal phase space of the witness bunch, shown in Fig. 3(a), highlights the effectiveness of using an increased witness bunch charge to flatten the wakefield throughout acceleration. The witness bunch is accelerated to a mean electron energy of 2.40 GeV, with a mean absolute deviation of 4.5 MeV (0.19%). The projected, normalized transverse emittance of the bunch is $\epsilon_{n,x/y} = 224$ nm in both the x - and y -planes. Figure 3(b) shows the longitudinal sliced parameters of the witness bunch, with a *slice-averaged* relative energy spread of 0.11% and transverse normalized emittance of $\epsilon_{n,\text{slice-av}} = 212$ nm in both planes.

Figures 3(c) and 3(d) demonstrate the evolution of $\epsilon_{n,x/y}$ and the trapped bunch charge, Q_b , throughout propagation. Here, an electron is defined to be trapped when its longitudinal momentum exceeds $p_z \geq 5$ MeV/c. Trapping first occurs at $z = 170$ μm , before the RPLP focal position ($z_{\text{foc}} = 0.2$ mm), with the trapped bunch charge increasing continuously until $z \approx 600$ μm . Figure 4(a) shows the evolution of the on-axis trapping potential, $\Delta\Psi = \Psi_i - \Psi_f$, throughout the first 1 mm of propagation. Here, $\Psi(\zeta) = -em_e^{-1}c^{-2} \int E_{z0}(\zeta)d\zeta$ represents the on-axis wakefield potential normalized by the electron rest mass; $\Psi_i = \Psi(\zeta \approx -45$ $\mu\text{m})$ represents the value of the on-axis wakefield potential at the location where electrons are initially released; and $\Psi_f = \Psi(\zeta = -87$ $\mu\text{m})$ corresponds to the rear of the cavity at $z = 0.2$ mm and approximate longitudinal location of electrons that become trapped. In this way, $|\Delta\Psi|$ represents the maximum normalized energy an electron can gain as it slips backward through the wakefield cavity—its value for each electron will be reduced according to the trajectory each takes as Ψ is maximized on axis. The trapping region is therefore represented by $\Delta\Psi \lesssim -1$ [37], and corresponds to the region for which an electron born at




FIG. 3. (a) Longitudinal phase space of the witness bunch after propagation over 90 mm. (b) Slice energy spread (gray, solid) and normalized transverse slice emittance (orange, dotted). In (b)–(d) the plotted quantities are separated by ionization source: He (blue); He^+ (pink); black lines represent calculations over the entire bunch. Evolution of the (c) projected normalized transverse emittance and (d) charge of the witness bunch over the entire propagation distance. In (c), thick solid (dashed) lines represent $\epsilon_{n,x}$ ($\epsilon_{n,y}$) for each ionization source; the thin black line represents $\langle \epsilon_n \rangle = \sqrt{\epsilon_{n,x}\epsilon_{n,y}}$ for the witness bunch. The vertical dotted line at $z = 1$ mm marks the handover between simulation codes.

rest within the wakefield gains enough longitudinal momentum to be trapped within the accelerating cavity. The boundaries of the on-axis trapping region are modified as electrons become trapped at the rear of the wakefield cavity, interfering with returning plasma electrons and briefly overloading the wakefield. This causes the shape of the wakefield potential to be modified between $z \approx 0.45$ and 0.6 mm [see Fig. 4(b)], reducing $\Delta\Psi$ for electrons released just prior to, and in, this region. This inhibits further trapping of electrons within the wakefield, as evidenced in Fig. 3(d). Beyond $z \approx 0.6$ mm, the RPLP intensity has dropped sufficiently that it no longer ionizes the dopant helium. Similarly, $\epsilon_{n,x/y}$ grows until $z = 0.6$ mm, but subsequently rapidly




FIG. 4. (a) Evolution of the on-axis trapping potential, $\Delta\Psi(\zeta, z) = \Psi_i - \Psi_f$, indicating the region for which released electrons can become trapped, $\Delta\Psi \lesssim -1$ (black, hatched). (b) Comparison of $\Delta\Psi$ just after trapping begins ($z = 0.2$ mm, black) and as the wakefield becomes overloaded ($z = 0.5$ mm, orange dashed).

decreases as initially trapped charge is lost as electrons at larger radii interact with the surrounding wakefield sheath and are accelerated away from the axis. After $z \approx 1.3$ mm both $\varepsilon_{n,x/y}$ and Q_b stabilize and remain constant thereafter, indicating matched propagation of the bunch.

Also shown in Figs. 3(b) and 3(c) are the sliced and projected (respectively) normalized transverse emittances of electrons originating from the two ionization levels of helium. The total bunch charge of $Q_b = 38.5$ pC is split, with 27.8 pC originating from the ionization of neutral He and 10.7 pC from ionization of He^+ . Interestingly, the contribution from He^+ , shown in pink, has a consistently lower $\varepsilon_{n,x/y}$ than the contribution from neutral He (blue).

Electrons released from the ionization of He^+ are born in regions of higher transverse—and also proportionally higher longitudinal—laser field than those released from neutral helium due to their higher ionization threshold. This increases the transverse momentum imparted to the electron by the laser field, resulting in a larger temperature—or thermal emittance, $\varepsilon_{n,th}$ —at ionization. However, within the wakefield the emittance of electrons released from ionization of He^+ is consistently smaller than for those from He. This is likely caused by the fact that He^+ is ionized in a region of smaller spatial extent. Figures 5(a) and 5(b) show the average normalized transverse momentum of ionized electrons, $\langle p_x(\zeta, x) \rangle$, both within (a) and without (b) the wakefield, in a single simulation timestep. In (a), initially the transverse momenta are dominated by the transverse laser field, but after passage of the RPLP the transverse wakefield becomes the dominant effect, increasing $\langle p_x \rangle$ by an order of magnitude with increasing delay behind the




FIG. 5. Evolution of the average transverse momentum, $\langle p_x \rangle$, of electrons ionized by the RPLP both (a) with and (b) without the presence of the wakefield. (c) Evolution of the transverse normalized slice emittance of electrons released from ionization of He (blue) and He^+ (pink) within a single simulation timestep, both with (solid) and without (dashed) the wakefield. Electrons slip backward (toward $-\zeta$) before being captured at the rear of the wakefield cavity ($\zeta - \zeta_{0,L} \approx -38$ μm). The longitudinal profile of the RPLP is represented by the gray shaded region. In (c), a median filter has been applied to the data to remove high frequency oscillations due to the laser field, whose effect can clearly be seen in (a) and (b). This timestep corresponds to $z = 0.2$ mm.

laser pulse, with momenta directed toward the axis ($p_x < 0$ for $x > 0$ and vice versa). Figure 5(c) shows the evolution of the transverse emittance: ahead of the peak of the laser pulse, $\zeta - \zeta_{0,L} > 0$, the growth of the emittance with decreasing ζ is similar with and without the wakefield. However, behind the peak the transverse wakefield begins to dominate, focusing electrons within the wakefield [Fig. 5(a)] and causing slight emittance growth as the electrons become matched to the transverse fields. Additional focusing of the electrons toward the axis by the transverse wakefield helps to combat their expansion from the buildup of the space-charge force, minimizing further emittance growth. When the electrons reach the rear of the wakefield cavity, $\zeta - \zeta_{0,L} \approx -38$ μm , the slice emittance grows significantly as electrons located off-axis interact with the returning plasma electron sheath.

III. COMPARISON TO THE LINEARLY-POLARIZED PLASMA PHOTOCATHODE

In order to isolate the effect of the polarization state, the simulation was repeated for a linearly-polarized (in y) laser pulse with all other parameters unchanged, including the




FIG. 6. (a) and (b) The transverse phase space of the trapped charge at (a) $z = 1$ mm and (b) $z = 90$ mm for the linearly-polarized (in y) case with laser parameters matching those of the RPLP in Figs. 1–3. The orange (purple) colorbars represent electrons released via ionization of He (He^+), and the contribution to the total charge from each ionization level at the corresponding z position is written above the plots. Comparison of the evolution of the (c) projected normalized transverse emittance, $\langle \epsilon_n \rangle = \sqrt{\epsilon_{n,x} \epsilon_{n,y}}$; (d) bunch charge; (e) relative energy spread; and (f) 6D brightness for the radially- (blue) and linearly-polarized (red) cases. The emittance components in x (pink, dot-dashed) and y (orange, dashed) are also shown for the linearly-polarized case in (c).

laser's peak intensity and transverse intensity profile. This meant both laser pulses ionized an equivalent volume of the HIT dopant, producing electron bunches of similar charge in both cases since the trapping dynamics are dominated by the beam-driven wakefield. Any further differences between the two cases could then be attributed to the effects of the ionizing pulse polarization alone. Figure 6(a) shows the transverse phase space of the trapped electrons following cessation of the injection process for the linearly-polarized pulse. In the $x-x'$ transverse phase space—the plane orthogonal to the polarization of the ionization pulse—the distributions of the trapped electrons are strongly peaked around the axis, closely mirroring the results for the RPLP [Fig. 2(a)]. However, in the plane of polarization of the ionization pulse ($y-y'$) the spatial distribution is far wider and peaks off axis. This broader distribution in $y-y'$ phase space is maintained throughout acceleration [Fig. 6(b)], leading to a larger projected emittance $\epsilon_{n,y} = 367$ nm for

the final electron bunch. In the x plane, the evolution of $\epsilon_{n,x}$ closely mirrors that of (both planes of) the RPLP case, with $\epsilon_{n,x} = 225$ nm at $z = 90$ mm, as shown in Fig. 6(d). The larger emittance in the y plane— $\epsilon_{n,y}/\epsilon_{n,x} \approx 1.6$ —is due to the contribution from the vector potential of the laser pulse at the moment of ionization, this time directed along y alone due to the linear polarization of the field. The fact that $\epsilon_{n,x}$ in the linearly-polarized case matches *both* planes of the RPLP case indicates that for this high-charge, optimally-loaded regime of the plasma photocathode, the emittance growth due to mismatch of the initial ionized electron distribution to the transverse wakefield dominates the final emittance of the accelerated electron bunch, rather than their initial thermal emittance. The initial emittance asymmetry results in different mismatch in each plane of the bunch, inducing an asymmetric emittance growth during the matching process, resulting in a larger geometric average emittance, $\langle \epsilon_n \rangle = \sqrt{\epsilon_{n,x} \epsilon_{n,y}}$, than for the RPLP case.

Simulations of the ionization of helium by both the radially- and linearly-polarized pulse *without* the presence of the wakefield further demonstrated the difference in thermal emittance for the two polarizations. The radially-polarized pulse produced electrons from ionization of neutral He with a thermal emittance $\epsilon_{n,\text{th}} \approx 33$ nm after passage of the laser pulse, and $\epsilon_{n,\text{th}} \approx 45$ nm for ionization of He^+ , as can be seen in Fig. 5(c). In both cases, the thermal emittances were symmetrical in x and y . In contrast, a strong asymmetry was observed for the linearly-polarized laser pulse with thermal emittances of $\epsilon_{n,\text{th},x} = 5$ nm (7 nm) and $\epsilon_{n,\text{th},y} = 50$ nm (75 nm) for electrons released via ionization of the first (second) level of helium. The corresponding evolution of the emittance within a single simulation timestep for the linearly-polarized laser pulse is shown in Fig. 7. Within the plane of polarization, the emittance evolution is similar both with and without the presence of the wakefield up until the




FIG. 7. Evolution of the slice transverse normalized emittance of electrons released from ionization of He and He^+ by the linearly-polarized (LP) laser pulse within a single simulation timestep, both with (solid) and without (dashed) the wakefield. The emittance in the plane of polarization, $\epsilon_{n,y}$, is shown in orange and represents electrons released via ionization of both He and He^+ . The same data are shown in red for $\epsilon_{n,x}$ (LP), and in gray for the RP pulse for comparison.

peak of the laser pulse, $\zeta - \zeta_{0,L} = 0$. After this point, $\zeta - \zeta_{0,L} < 0$, the mismatch with the transverse wakefield leads to slight emittance growth in $\varepsilon_{n,y}$ relative to when there is no wakefield present. The maximum $\varepsilon_{n,y}$, occurring near the location of the peak intensity of the laser pulse as the electrons undergo quiver motion in the laser field, is larger than the equivalent for the RPLP [Fig. 5(c)] due to the increased transverse field amplitude of the linearly-polarized pulse—for equal intensity, $E_{\max}(\text{LP}) \simeq \sqrt{2}E_{\max}(\text{RP})$. Out of the plane of polarization ($\varepsilon_{n,x}$), the difference with and without the presence of the wakefield is much more marked, highlighting the effect of the transverse mismatch and the severity of the emittance growth it induces. In the presence of the wakefield, the magnitude of the emittance asymmetry ($\varepsilon_{n,y}/\varepsilon_{n,x}$) is much reduced after passage of the laser pulse, but the larger emittance in the plane of polarization remains.

The witness bunch charge is approximately equal for both polarizations, with $Q_b = 38.8$ pC for the linearly-polarized ionization pulse [Fig. 6(d)], confirming that the longitudinal field present at the focus of the RPLP has minimal impact on the electron trapping process. The energy spectrum remains similar to the RPLP case, with the witness bunch reaching a mean energy of 2.38 GeV and a mean absolute deviation of 8.0 MeV (0.33%), slightly larger than for the RPLP. Figure 6(e) shows the evolution of both projected and slice-averaged relative energy spread. The slice-averaged spread decreases with propagation for both polarization states, indicating effective local flattening of the accelerating field by the bunch core, which limits growth in absolute slice spread as energy increases. This suggests the higher projected energy spread in the LP case arises from particles in the tails, outside the optimally-loaded region. Figure 6(f) compares the 6D brightness evolution, highlighting the advantage of the RPLP-generated bunch: its projected brightness, $B_{\text{6D,RP}}^{(\text{p})} = 5.9 \times 10^{25} \text{ A m}^{-3}$, is 2.4 times higher than the LP case ($2.4 \times 10^{25} \text{ A m}^{-3}$), and its slice-averaged brightness, $B_{\text{6D,RP}}^{(\text{s})} = 1.1 \times 10^{26} \text{ A m}^{-3}$, exceeds the LP value ($9.3 \times 10^{25} \text{ A m}^{-3}$) by a factor of 1.2.

IV. MULTIOBJECTIVE BAYESIAN OPTIMIZATION

In order to better understand the performance of the radially-polarized plasma photocathode scheme, a multi-objective Bayesian optimization (MOBO) [38,39] routine was used to explore the input parameter space. MOBO aims to find optimal trade-offs between multiple, often conflicting objectives to define the Pareto front (PF) of the system—the region of the parameter space for which one objective cannot be improved without deleterious effects on another. During performance of the MOBO routine, a surrogate Gaussian process (GP) model [40] is continually updated to enable cheap, noise-free predictions of the outcome of simulations to be made and thus generate a search strategy

to identify promising regions of the input parameter space to be explored. Here, the routine uses (noisy) expected hyper-volume improvement [41] as the acquisition function.

Four input parameters were varied: the total energy of the ionization laser pulse, E_{Laser} ; its spot size at focus, w_0 ; the fraction of helium dopant, n_{dop} ; and the relative delay between the centroid of the drive beam and ionization laser pulse, $\Delta\zeta$. The range over which these parameters were allowed to vary is summarized in Table I. The first three of these parameters were used to control the quantity of charge released into the wakefield via ionization of helium, while variation of the relative delay, $\Delta\zeta$, can be used to study the effects of timing fluctuations on the injection process.

The MOBO routine, performed using the Optimas [42,43] Python library, aimed to simultaneously optimize three key metrics of the accelerated electron bunch: its charge, Q_b ; projected relative energy spread, $\sigma_{E,\text{rel}} = \sigma_{E,\text{mad}}/E_{\text{mean}}$, where $\sigma_{E,\text{mad}}$ is the rms absolute deviation from the mean, E_{mean} ; and its geometric average transverse normalized emittance, $\langle \varepsilon_n \rangle = \sqrt{\varepsilon_{n,x} \cdot \varepsilon_{n,y}}$. The value of an “overall” optimization metric,

$$\mathcal{F} = -\ln \left[\frac{E_{\text{mean}}}{1 \text{ GeV}} \left(\frac{Q_b}{10 \text{ pC}} \right)^{1/2} \cdot \left(\frac{\sigma_{E,\text{rel}}}{0.5\%} \right)^{-1} \left(\frac{\langle \varepsilon_n \rangle}{100 \text{ nm}} \right)^{-1} \right], \quad (2)$$

which combines the three key output parameters in a way such that minimizing $\mathcal{F}$ represents electron bunches with improved quality was also tracked, along with additional parameters such as the relative contributions of electrons released from the different ionization levels of helium to the final bunch charge.

In total, 500 simulations for the RPLP setup were performed, with a further 200 simulations optimizing the linearly-polarized case for comparison (see the Appendix for details of the start-to-end modeling workflow). Figure 8 demonstrates the predictions of the resulting surrogate GP model after 500 simulations of the RPLP setup. Two contour plots are shown, representing the predicted value of $\mathcal{F}$ in 2D slices through the four-dimensional input parameter space. The slices are taken at locations close to input parameters that generated the highest quality witness bunch, defined by minimizing the value of $\mathcal{F}$:

TABLE I. Summary of parameters used within the RPLP MOBO study and their minimum, maximum, and optimum values.

Parameter	Minimum	Maximum	Optimum
E_{Laser} (mJ)	0.05	1.0	0.3
w_0 (μm)	2.0	10.0	4.0
n_{dop} (%)	1.0	10.0	5.0
$\Delta\zeta$ (μm)	40.0	60.0	48.0




FIG. 8. Visualization of the GP model fitted to the results of the MOBO routine. Two slices along the four-dimensional input parameter space are shown, taken at values close to those at which the minimum value of $\mathcal{F}$ was achieved: $E_{\text{Laser}} = 0.3$ mJ, $w_0 = 4.0$ μm , $n_{\text{dop}} = 5\%$, and $\Delta\zeta = 48$ μm . Contour plots of the value of $\mathcal{F}$ predicted by the fitted GP model are shown as a function of (a) w_0 and E_{Laser} , and (b) $\Delta\zeta$ and n_{dop} . For each plot, GP model predictions and uncertainties are shown for the slices along the optimal input parameters.

$E_{\text{Laser}} = 0.3$ mJ, $w_0 = 4.0$ μm , $n_{\text{dop}} = 5\%$, and $\Delta\zeta = 48$ μm . While visualization of the GP model in this way allows qualitative understanding of the stability of the input parameter space surrounding the optimal parameters, use of a single value, $\mathcal{F}$, to represent the quality of the outcome does not permit easy interpretation of correlations with individual witness bunch parameters.

A. Stability

To better quantify the stability of the input parameter space around the optimum, Monte-Carlo (MC) sampling ($N_{\text{samples}} = 10^4$) of the surrogate model was performed. Typical experimental fluctuations were added to the optimal input parameters, sampled according to a normal distribution centered on the optimal value with a chosen relative width. The relative magnitude of the fluctuations were: $E_{\text{Laser}} = 3\%$, $w_0 = 10\%$, $n_{\text{dop}} = 10\%$, and $\Delta\zeta = 4\%$, chosen to represent typical fluctuations observed in experiments. Initially, each input parameter was varied individually, while keeping the others fixed to isolate the contributions to the fluctuations of the witness parameters. The dominant source of fluctuations was found to be the relative delay between the drive bunch and ionization laser pulse, $\Delta\zeta$, with a magnitude corresponding to a timing jitter of ± 10 fs, similar to state-of-the-art synchronization systems [44–46]. The resulting witness bunch parameters varied over the range $Q_b = (27.3 \pm 6.7)$ pC, $\langle \epsilon_n \rangle = (256 \pm 33)$ nm, and $\sigma_{E,\text{rel}} = (0.56 \pm 0.48)\%$, where the

quoted fluctuations represent the standard deviation. The variations of the key witness bunch parameters resulting from the MC sampling are shown in Fig. 9. Also shown are the results of PIC simulations (red data points) performed using the optimal input parameters but varying only $\Delta\zeta$ over the same range as tested with the MC sampling. The two agree well, confirming the quality of the surrogate model and its ability to make accurate predictions.

The witness bunch charge can be seen to be strongly correlated with the relative delay of the ionization laser pulse [Fig. 9(a)], with an increased bunch charge at smaller delays. This is because the ionization pulse centroid moves




FIG. 9. Distribution of witness bunch parameters resulting from MC sampling of the surrogate GP model when an rms timing fluctuation of ± 10 fs around the optimum is introduced. Red data points represent results from PIC simulations when varying $\Delta\zeta$ but keeping all other input parameters fixed. $\zeta_{0,L}$ is the location of the RPLP centroid. (a) Q_b [pC], (b) $\langle \epsilon_n \rangle$ [nm], and (c) $\sigma_{E,\text{rel}}$ [%].

closer to the wakefield potential minimum— $\zeta \approx -38 \mu\text{m}$ (Fig. 4)—maximizing $|\Delta\Psi|$ for the released electrons. However, this has two subsequent effects: the transverse emittance of the bunch increases and so does its relative energy spread as demonstrated in Figs. 9(b) and 9(c), respectively. The increase in bunch charge causes the wakefield to become locally overloaded, causing an energy chirp to be imparted to the witness bunch during acceleration. Conversely, for larger delays, electrons are released further back in the trapping region [Fig. 4(b)] and a smaller proportion of electrons gain enough longitudinal momentum to become trapped. This reduces the bunch charge and emittance, but now underloads the wakefield, again causing an increased energy spread.

Such sensitivity to the relative timing between the drive bunch and laser pulse makes this scheme experimentally challenging; however, there are ways in which this can be overcome. For example, operating the accelerator at a lower plasma density increases the size of the wakefield cavity and significantly relaxes the timing constraints. At the on-axis density used in this study, $n_0 = 1 \times 10^{17} \text{cm}^{-3}$, a timing fluctuation of ± 10 fs represents approximately 6% of the duration of the wakefield cavity. Reducing the on-axis density to that proposed for future plasma-based accelerator facilities [34], $n_0 = 7 \times 10^{15} \text{cm}^{-3}$, increases the length of the trapping region such that the same fluctuation represents less than 2% of the cavity, which would reduce the impact of such fluctuations. Furthermore, finding a combination of input parameters that enables optimal beam loading, where the position of the ionizing laser pulse corresponds to the minimum of the wakefield potential could further reduce the timing sensitivity thanks to the parabolic shape of the wakefield potential around its minimum [5]. We note also that recent work has demonstrated that active energy compression can reduce the shot-to-shot jitter of laser-accelerated electron bunches to the per mille level [47].

We note that the relative fluctuations in witness bunch parameters due to temporal jitter found in this study exceed those reported in Ref. [5], which were at most 5% ($\sigma_{E,\text{rel}}$), if not significantly less. However, the scenario modeled in Ref. [5] focused on minimizing witness bunch emittance by using pulse intensities just above the ionization threshold of the dopant species, resulting in lower bunch charges (~ 1 pC). Furthermore, in that work, the simulations considered witness acceleration over only 8 mm, and not until drive bunch depletion. As discussed previously, the lower bunch charge precludes loading of the wakefield and causes the energy spread of the witness bunch to grow significantly during propagation, which will result in larger energy fluctuations by the end of the accelerator. To provide a point of comparison with this earlier work, the RPLP simulation results obtained in this work were sampled in a region of the parameter space that provided low-charge and ultralow emittance witness bunches—achieved by minimizing the ionizing laser pulse energy and spot size. The MC procedure described above yielded, for this working

point, witness bunch parameters of $Q_b = (2.8 \pm 1.4)$ pC, $\langle \epsilon_n \rangle = (64 \pm 11)$ nm, and $\sigma_{E,\text{rel}} = (2.5 \pm 0.6)\%$ when subject to the same temporal jitter as the optimally-loaded working point. The transverse emittance values are competitive with those found in previous studies [5,12–15], despite the increased thermal emittance in both transverse planes introduced by the field of the RPLP. However, the large relative energy spread and energy jitter obtained for the low-charge case would be challenging for further beam transport. The high-charge, optimally beam-loaded regime identified in this study would reduce problems of this kind.

Further fluctuations in witness bunch parameters will be induced by fluctuations of the transverse focal position of the ionizing laser pulse. Using a short focal length optic to tightly focus the beam reduces such jitter and remaining fluctuations can be further suppressed by recent beam stabilization advances [48]. Scaling the results of Amodio *et al.* to our case suggests that the jitter in the transverse position of the injection pulse could be of order $1 \mu\text{m}$, which is small compared to the transverse size of the plasma cavity. Modeling of this effect would be an interesting topic for future work, but would require either 3D simulations of the interaction or the use of analytic pulse techniques [49] in quasi-3D geometries.

We note also that the use of a higher-charge “escort” bunch to locally flatten the longitudinal wakefield in the region of the low-charge witness [15] not only significantly increases experimental complexity but its generation will likely also be subject to similar fluctuations as those found in this study, inducing further fluctuations of the witness bunch parameters.

B. Performance

Further understanding of the ultimate performance of the radially-polarized plasma photocathode can be gleaned by exploring the relationship between bunch charge and emittance. Maximizing bunch charge and consequently peak current, while minimizing transverse emittance is crucial for maximizing the 6D brightness of the electron bunch. Generation of a high-brightness electron bunch using a traditional photocathode is in part governed by a fundamental limitation relating the charge of the bunch to its emittance as a result of space-charge forces, an effect that is particularly dominant before electrons gain significant longitudinal momentum. A similar relation can be extracted for the plasma photocathode from the results of the MOBO routine. An example of this is presented in Fig. 10(a), which shows the relation between the final projected normalized average transverse emittance of the accelerated witness bunch, $\langle \epsilon_n \rangle$, and its charge, Q_b , for a radially-polarized ionization pulse. The Pareto front, indicated by the dashed blue line, represents the minimum achievable emittance of the witness bunch as a function of the bunch charge. A power law can be fitted to extract the relationship between the two: a fit performed using Bayesian linear regression indicated a power of (0.48 ± 0.03) , where the uncertainty




FIG. 10. (a) Relation between Q_b and $\langle \epsilon_n \rangle$ for the RPLP. The size of data points represents their value of $\mathcal{F}$, with larger points representing higher quality witness bunches. The color of a data point indicates the fraction of the final witness bunch charge from ionization of He^+ . The extracted PF is demonstrated by the dashed, blue line. (b) Comparison between the PFs for radial (blue) and linear polarization (red). The components of ϵ_n in the x - (pink, dot-dashed) and y -planes (orange, dashed) are also shown for linear polarization. (c) $\epsilon_{n,y}/\epsilon_{n,x}$ as a function of ionization pulse normalized vector potential, a_0 . (d) and (e) Mean (line) and standard error (shaded region) of predictions from the surrogate GP model demonstrating the variation in $\langle \epsilon_n \rangle$ with (d) w_0 and (e) E_{Laser} while keeping all other parameters fixed at their optima. The dashed gray vertical lines represent the optimal values of the laser parameters, as defined in Table I.

represents the 95% confidence interval, therefore consistent with a scaling of $\langle \epsilon_n \rangle \propto \sqrt{Q_b}$ in this regime. While a separate study would be required for alternative drive bunch parameters, this scaling indicates μm -scale transverse emittances could be achieved with the radially-polarized plasma photocathode scheme when generating nC-class bunches, matching those required by future plasma-based linear collider designs [34].

Figure 10(b) compares the extracted PFs for the radially- and linearly-polarized cases. The PF is very similar for both, indicating that in this regime the final emittance of the witness bunch is dominated by its evolution during the trapping and matching phase—defined by the properties of the wakefield and therefore consistent in both cases—and not by its thermal emittance at the point of ionization. Also shown are the corresponding PFs for the x - and y -components of the emittance in the linearly-polarized case, demonstrating its asymmetry. Figure 10(c) shows the emittance asymmetry at $z = 90$ mm as a function of the ionization laser pulse normalized vector potential, $a_0 = (e\lambda E_{\text{max}})/(2\pi m_e c^2)$, where E_{max} is the maximum of the laser pulse field. While the emittance remains symmetric over the entire range of a_0 simulated in the radially-polarized case (blue), the asymmetry grows with increasing a_0 for the linearly-polarized case (red), illustrating a key advantage of the radially-polarized scheme with the increased laser intensity necessary for high-charge operation.

It is important to note that while the PFs are very similar for both polarization states, the input parameters required to achieve minimal emittance can differ. When keeping the laser pulse parameters fixed and only changing its polarization, the average transverse emittance of the bunch is reduced with a RPLP when operating in the regime of optimal beam loading of the wakefield, as shown previously in Fig. 6. Figures 10(d) and 10(e) demonstrate this effect over the entire range of ionizing laser pulse parameters considered. The figure shows the variation of the witness bunch emittance with (d) spot size and (e) pulse energy for the radially- and linearly-polarized ionization pulse while keeping all other parameters fixed at their optimal values. In the region of the parameter space around which the witness bunch optimally loads the wakefield and maximizes its 6D brightness, indicated by the dashed gray lines, the average transverse emittance of the witness bunch is improved when using a radially-polarized ionizing laser pulse compared to linear polarization.

At the optimal input parameters found using the MOBO routine, the resulting witness bunch had a projected 6D brightness of $B_{6D}^{(p)} = 2.0 \times 10^{26} \text{ A m}^{-3}$, while the slice-averaged brightness peaked at $B_{6D}^{(s)} = 3.8 \times 10^{26} \text{ A m}^{-3}$. For effective FEL generation, the dimensionless Pierce parameter ρ is constrained by the relative energy spread of the bunch [50], i.e., $\rho \geq \Delta E/E$. Here, the slice-averaged energy spread is found to be $\sigma_{E,\text{rel,sl}} = 1.1 \times 10^{-3}$ with a minimum relative slice energy spread of $\approx 4 \times 10^{-4}$, matching requirements of high-gain FELs. The minimum radiation wavelength produced by such a bunch can be calculated in the 1D limit via $\lambda_r \geq 4\pi\epsilon_n/\gamma_0$ [51], where γ_0 is the relativistic Lorentz factor of the bunch. The ultralow emittance and high energies of the bunches produced in this scheme result in a minimum radiation wavelength of $\lambda_r \geq 0.6$ nm. Lasing at shorter wavelengths could be achieved by using a higher energy drive bunch to increase the energy of the witness bunch before drive depletion. A key advantage of operation

in the optimally-loaded regime of the plasma photocathode is that increasing the acceleration length allows increase of the mean energy of the witness bunch without degradation of the energy spread or transverse emittance, thanks to the flattened longitudinal field gradient and matched propagation. The brightness values achieved in this study match, and even exceed, those currently produced at state-of-the-art XFEL facilities [6], and demonstrate the ability of this scheme to transform the brightness of existing light sources with minimal impact on the footprint of the facility. For example, the drive bunch parameters chosen in this study closely match those provided by the FLASH facility at DESY [52] and could increase its 6D brightness by multiple orders of magnitude through the implementation of this scheme, while pushing the radiation wavelength below the nm-scale—currently FLASH I/II operates at $B_{6D}^{(p)} \sim \mathcal{O}(10^{22} \text{ A m}^{-3})$, $B_{6D}^{(s)} \sim \mathcal{O}(10^{23} \text{ A m}^{-3})$, $\lambda_r = 4\text{--}6 \text{ nm}$.

V. CONCLUSIONS

Using start-to-end simulations we have demonstrated the high-charge regime of plasma photocathode operation, enabling optimal beam loading from a single injection event to simultaneously minimize the projected energy spread and maximize the 6D brightness of the witness bunch. Use of a radially-polarized ionizing laser pulse enables generation of witness bunches with symmetric emittance in the x - and y -planes, improving the 6D brightness of the witness bunch compared to an equivalent linearly-polarized ionizing laser pulse in the optimally-loaded regime. Furthermore, having identical emittance in both transverse planes eliminates the risk of emittance mixing during acceleration and eases beam transport and matching after the accelerator, while offering improved radiation generation opportunities, for example, in helical undulators.

In the high-charge regime of the plasma photocathode scheme, trapping of electrons released from higher ionization levels of the dopant species was shown to reduce the emittance of the bunch due to the reduced spatial extent of the ionized electrons, in contrast to behavior outside of the wakefield environment. This is promising for controlled electron injection schemes relying on localized, selective ionization within quasilinear laser-wakefield accelerator (LWFA) stages [53,54]. In particular, it suggests the plasma photocathode scheme could be used to generate high-quality electron bunches within LWFA stages driven by a train of laser pulses [55,56] as the reduced peak vector potential ($a_0 \ll 1$) relative to single-pulse LWFA could still allow operation at low ionization laser pulse energies.

Multiobjective Bayesian optimization was used to explore the stability and performance of the high-charge regime, indicating particular sensitivity to the relative timing of the drive bunch and ionizing laser pulse. Extraction of the Pareto front relating the bunch charge to its average transverse normalized emittance indicated a scaling $\langle \epsilon_n \rangle \propto \sqrt{Q_b}$,

suggesting the generation of μm -scale emittances for nC-scale bunch charges as required for future plasma-based facilities. This study introduces a framework [57] for further exploration of the (radially-polarized) plasma photocathode scheme where optimization can be performed with specific use cases in mind. Inclusion of additional experimental parameters such as plasma length, density, and (dopant) species may be appropriate in this case.

ACKNOWLEDGMENTS

The authors would like to acknowledge useful discussions with R. D’Arcy. This work was supported by the UK Science and Technology Facilities Council (STFC UK), Grants No. ST/V001655/1 and Agreement No. 9253 S2 2021 010; the Engineering and Physical Sciences Research Council, Grants No. EP/R513295/1 and No. EP/V006797/1. Computing resources provided by STFC Scientific Computing Department’s SCARF cluster. This research was funded in whole, or in part, by EPSRC and STFC, which are Plan S funders.

DATA AVAILABILITY

The data that support the findings of this article are openly available [57].

APPENDIX: START-TO-END MODELING

Modeling of the first 1 mm of the interaction was performed with the FBPIC (v.0.26.0) particle-in-cell code using $N_m = 3$ azimuthal modes. The injection simulations corresponding to results in Figs. 1–7 were performed in the lab frame, with 48 particles per cell (PPC) for each species: preionized plasma electrons and the dopant neutral helium. A nonevolving background of ions was assumed as there was no discernible difference in wakefield or witness bunch parameters seen when including mobile hydrogen ions. The simulation domain measured $60 \times 157 \mu\text{m}$ ($r \times z$), with $\Delta r = 0.5 \mu\text{m}$ and $\Delta z = 0.02 \mu\text{m}$ (i.e., $\lambda/40$). The injected witness bunch was typically formed of $\sim 4 \times 10^6$ macroparticles (depending on its charge), while the Gaussian drive bunch was modeled using 1×10^6 equal weight macroparticles.

At the end of the ionization simulation, the drive and witness bunch distributions were written to file and then loaded into the extended propagation simulation in Wake-T (v.0.8.0). The witness bunch distributions were separated according to the ionization source, either neutral He or He^+ , to enable tracking of the separate components. The Wake-T simulation used an identical size of simulation domain with $\Delta r = 0.25 \mu\text{m}$ and $\Delta z = 0.14 \mu\text{m}$.

The propagation distance at which the handover between the ionization (FBPIC) and extended propagation (Wake-T) simulations was performed was scanned between $z_h = 0.6$ and 2.0 mm . With $z_h = 0.6 \text{ mm}$, the final witness bunch charge after drive bunch depletion at $z = 90 \text{ mm}$ was lower due to electrons released from ionization of

helium later in the simulation not yet meeting the “trapped” criterion ($p_z \geq 5$ MeV/c) by the handover position, z_h . For handover positions $z_h \geq 1.0$ mm the witness bunch charge (and all other properties) was consistent at $z = 90$ mm in all cases, and therefore a handover position of $z_h = 1.0$ mm was selected to minimize the duration of the ionization simulation.

The extended propagation simulation performed using Wake-T was first benchmarked against a full 3D HiPACE++ (v.24.03) [58] simulation using the same drive and witness bunches output by a test FBPIC simulation after 1 mm of propagation. The benchmark HiPACE++ simulation had a global grid resolution of $\Delta x = \Delta y = 0.26$ μm , $\Delta z = 0.08$ μm with 4 PPC for plasma electrons and 1 PPC for (mobile) ions. Around the witness bunch, a region of radius 12.5 μm was further resolved using mesh refinement, with an increased resolution of $\Delta x = \Delta y = 0.032$ μm with 256 PPC for plasma electrons and 64 PPC for ions. For both the Wake-T and HiPACE++ simulations, an adaptive timestep (Δt) was used with value corresponding to $\omega_\beta \Delta t = 2\pi/20$, where $\omega_\beta = \omega_{p0}/\sqrt{2\gamma}$ with ω_{p0} the plasma angular frequency and γ an average of Lorentz factors of the slowest particles in both the drive and witness beams. The benchmark HiPACE++ simulation took approximately 200 min to perform using 4 $\times$ NVIDIA A100 (40 GB) GPUs, while the Wake-T simulation took approximately 20 min on a single AMD EPYC 7502 CPU core. A comparison of the witness bunch parameters for the two simulations after 90 mm propagation are shown in Table II and show excellent agreement, justifying the use of Wake-T for extended propagation simulations thanks to the reduced computational requirements.

Ionization simulations performed with FBPIC during the MOBO routine (Figs. 8–10) used the boosted frame ($\gamma_{\text{boost}} = 2$) to reduce computational time. Convergence tests demonstrated consistent witness bunch parameters between the lab and boosted frames, although with fewer macroparticles representing the witness bunch in the boosted frame—on average, the bunch was represented by $\sim 5 \times 10^5$ macroparticles. The size of the simulation domain, resolution used and PPC matched those of the lab frame simulation, although these simulations also included mobile plasma ions.

TABLE II. Comparison of witness bunch parameters from a test simulation after 90 mm propagation using HiPACE++ (including mesh refinement and ion motion) and Wake-T.

Parameter	HiPACE++	Wake-T
E_{mean} (GeV)	2.56	2.57
$\sigma_{E,\text{rel}}$ (%)	3.8	3.9
I_{pk} (kA)	3.1	3.1
Q_b (pC)	17.7	17.6
$\epsilon_{n,x}$ (nm)	219	221
$\epsilon_{n,y}$ (nm)	220	222

Handover between the various steps of the simulation workflow was handled by the ChainEvaluator function available within the Optimas (v.0.6.0) Python package. For each trial (i.e., set of input variables), a single NVIDIA A100 (40 GB) GPU was used to perform the ionization simulation in FBPIC, followed by the extended propagation simulation in Wake-T performed on a single AMD EPYC 7502 CPU core. A single start-to-end simulation took approximately 100 min to complete, with 80 min required for injection modeling and 20 min for extended propagation. Four separate simulations could be run in parallel, with the entire optimization routine performed in approximately 200 h.

The version of the LASY Python package used in this study was v.0.4.0. Simulation input scripts for the results presented in this study can be found at [57].

-
- [1] C. Pellegrini, A. Marinelli, and S. Reiche, The physics of x-ray free-electron lasers, *Rev. Mod. Phys.* **88**, 015006 (2016).
 - [2] E. A. Seddon, J. A. Clarke, D. J. Dunning, C. Masciovecchio, C. J. Milne, F. Parmigiani, D. Rugg, J. C. H. Spence, N. R. Thompson, K. Ueda, S. M. Vinko, J. S. Wark, and W. Wurth, Short-wavelength free-electron laser sources and science: A review, *Rep. Prog. Phys.* **80**, 115901 (2017).
 - [3] I. Tetsuya, Accelerator-based x-ray sources: Synchrotron radiation, x-ray free electron lasers and beyond, *Phil. Trans. R. Soc. A* **377**, 20180231 (2019).
 - [4] S. Shin, New era of synchrotron radiation: Fourth-generation storage ring, *AAPPS Bull.* **31**, 21 (2021).
 - [5] A. F. Habib *et al.*, Plasma photocathodes, *Ann. Phys. (Amsterdam)* **535**, 2200655 (2023).
 - [6] S. Di Mitri and M. Cornacchia, Electron beam brightness in linac drivers for free-electron-lasers, *Phys. Rep.* **539**, 1 (2014).
 - [7] T. Tajima and J. M. Dawson, Laser electron accelerator, *Phys. Rev. Lett.* **43**, 267 (1979).
 - [8] R. D. Ruth, A. W. Chao, P. L. Morton, and P. B. Wilson, A plasma wake field accelerator, Technical Report No. SLAC-PUB-3374, SLAC National Accelerator Laboratory (SLAC), Menlo Park, CA, 1984.
 - [9] P. Chen, J. M. Dawson, R. W. Huff, and T. Katsouleas, Acceleration of electrons by the interaction of a bunched electron beam with a plasma, *Phys. Rev. Lett.* **54**, 693 (1985).
 - [10] P. Chen, S. Cheshkov, R. Ruth, and T. Tajima, An ultra-high gradient plasma wakefield booster, *AIP Conf. Proc.* **569**, 903 (2001).
 - [11] I. Blumenfeld, C. E. Clayton, F.-J. Decker, M. J. Hogan, C. Huang, R. Ischebeck, R. Iverson, C. Joshi, T. Katsouleas, N. Kirby, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, E. Oz, R. H. Siemann, D. Walz, and M. Zhou, Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator, *Nature (London)* **445**, 741 (2007).
 - [12] B. Hidding, G. Pretzler, J. B. Rosenzweig, T. Königstein, D. Schiller, and D. L. Bruhwiler, Ultracold electron bunch

- generation via plasma photocathode emission and acceleration in a beam-driven plasma blowout, *Phys. Rev. Lett.* **108**, 035001 (2012).
- [13] C. B. Schroeder, J.-L. Vay, E. Esarey, S. S. Bulanov, C. Benedetti, L.-L. Yu, M. Chen, C. G. R. Geddes, and W. P. Leemans, Thermal emittance from ionization-induced trapping in plasma accelerators, *Phys. Rev. ST Accel. Beams* **17**, 101301 (2014).
- [14] B. Hidding, G. G. Manahan, O. Karger, A. Knetsch, G. Wittig, D. A. Jaroszynski, Z.-M. Sheng, Y. Xi, A. Deng, J. B. Rosenzweig, G. Andonian, A. Murokh, G. Pretzler, D. L. Bruhwiler, and J. Smith, Ultrahigh brightness bunches from hybrid plasma accelerators as drivers of 5th generation light sources, *J. Phys. B* **47**, 234010 (2014).
- [15] A. F. Habib, G. G. Manahan, P. Scherkl, T. Heinemann, A. Sutherland, R. Altuiri, B. M. Alotaibi, M. Litos, J. Cary, T. Raubenheimer, E. Hemsing, M. J. Hogan, J. B. Rosenzweig, P. H. Williams, B. W. J. McNeil, and B. Hidding, Attosecond-angstrom free-electron-laser towards the cold beam limit, *Nat. Commun.* **14**, 1054 (2023).
- [16] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva, Beam loading in the nonlinear regime of plasma-based acceleration, *Phys. Rev. Lett.* **101**, 145002 (2008).
- [17] S. Diederichs, C. Benedetti, A. F. Pousa, A. Sinn, J. Osterhoff, C. B. Schroeder, and M. Thévenet, Resonant emittance mixing of flat beams in plasma accelerators, *Phys. Rev. Lett.* **133**, 265003 (2024).
- [18] S. C. Tidwell, D. H. Ford, and W. D. Kimura, Generating radially polarized beams interferometrically, *Appl. Opt.* **29**, 2234 (1990).
- [19] S. C. Tidwell, G. H. Kim, and W. D. Kimura, Efficient radially polarized laser beam generation with a double interferometer, *Appl. Opt.* **32**, 5222 (1993).
- [20] R. Dorn, S. Quabis, and G. Leuchs, Sharper focus for a radially polarized light beam, *Phys. Rev. Lett.* **91**, 233901 (2003).
- [21] G. Machavariani, Y. Lumer, I. Moshe, A. Meir, and S. Jackel, Efficient extracavity generation of radially and azimuthally polarized beams, *Opt. Lett.* **32**, 1468 (2007).
- [22] N. Zaïm, D. Guénot, L. Chopineau, A. Denoeud, O. Lundh, H. Vincenti, F. Quéré, and J. Faure, Interaction of ultraintense radially-polarized laser pulses with plasma mirrors, *Phys. Rev. X* **10**, 041064 (2020).
- [23] J. Powell, S. W. Jolly, S. Vallières, F. Fillion-Gourdeau, S. Payeur, S. Fourmaux, M. Lytova, M. Piché, H. Ibrahim, S. MacLean, and F. m. c. Légaré, Relativistic electrons from vacuum laser acceleration using tightly focused radially polarized beams, *Phys. Rev. Lett.* **133**, 155001 (2024).
- [24] S. W. Jolly and M. A. Porras, Clarification for the fields of different radially polarized Laguerre–Gaussian light beams, *Opt. Lett.* **47**, 3632 (2022).
- [25] M. Thévenet, I. A. Andriyash, L. Fedeli, A. F. Pousa, A. Huebl, S. Jalas, M. Kirchen, R. Lehe, R. J. Shalloo, A. Sinn, and J.-L. Vay, LASY: LAsEr manipulations made eaSY, *arXiv:2403.12191*.
- [26] R. Lehe, M. Kirchen, I. A. Andriyash, B. B. Godfrey, and J.-L. Vay, A spectral, quasi-cylindrical and dispersion-free particle-in-cell algorithm, *Comput. Phys. Commun.* **203**, 66 (2016).
- [27] S. Jalas, I. Dornmair, R. Lehe, H. Vincenti, J.-L. Vay, M. Kirchen, and A. R. Maier, Accurate modeling of plasma acceleration with arbitrary order pseudo-spectral particle-in-cell methods, *Phys. Plasmas* **24**, 033115 (2017).
- [28] M. Kirchen, R. Lehe, S. Jalas, O. Shapoval, J.-L. Vay, and A. R. Maier, Scalable spectral solver in Galilean coordinates for eliminating the numerical Cherenkov instability in particle-in-cell simulations of streaming plasmas, *Phys. Rev. E* **102**, 013202 (2020).
- [29] A. Ferran Pousa, R. Assmann, and A. Martinez de la Ossa, Wake-T: A fast particle tracking code for plasma-based accelerators, *J. Phys. Conf. Ser.* **1350**, 012056 (2019).
- [30] R. D’Arcy *et al.*, FLASHForward: Plasma wakefield accelerator science for high-average-power applications, *Phil. Trans. R. Soc. A* **377**, 20180392 (2019).
- [31] V. Yakimenko, L. Alsberg, E. Bong, G. Bouchard, C. Clarke, C. Emma, S. Green, C. Hast, M. J. Hogan, J. Seabury, N. Lipkowitz, B. O’Shea, D. Storey, G. White, and G. Yocky, FACET-II facility for advanced accelerator experimental tests, *Phys. Rev. Accel. Beams* **22**, 101301 (2019).
- [32] T. Kurz *et al.*, Demonstration of a compact plasma accelerator powered by laser-accelerated electron beams, *Nat. Commun.* **12**, 2895 (2021).
- [33] F. M. Foerster *et al.*, Stable and high-quality electron beams from staged laser and plasma wakefield accelerators, *Phys. Rev. X* **12**, 041016 (2022).
- [34] B. Foster, R. D’Arcy, and C. A. Lindstrøm, A hybrid, asymmetric, linear Higgs factory based on plasma-wakefield and radio-frequency acceleration, *New J. Phys.* **25**, 093037 (2023).
- [35] S. Kawata, Q. Kong, S. Miyazaki *et al.*, Electron bunch acceleration and trapping by ponderomotive force of an intense short-pulse laser, *Laser Part. Beams* **23**, 61 (2005).
- [36] D. N. Gupta, N. Kant, D. E. Kim, and H. Suk, Electron acceleration to GeV energy by a radially polarized laser, *Phys. Lett. A* **368**, 402 (2007).
- [37] A. Pak, K. A. Marsh, S. F. Martins, W. Lu, W. B. Mori, and C. Joshi, Injection and trapping of tunnel-ionized electrons into laser-produced wakes, *Phys. Rev. Lett.* **104**, 025003 (2010).
- [38] M. Emmerich, K. Giannakoglou, and B. Naujoks, Single- and multiobjective evolutionary optimization assisted by Gaussian random field metamodels, *IEEE Trans. Evol. Comput.* **10**, 421 (2006).
- [39] S. Daulton, M. Balandat, and E. Bakshy, Differentiable expected hypervolume improvement for parallel multi-objective Bayesian optimization, in *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS ’20* (Curran Associates Inc., Red Hook, NY, 2020).
- [40] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, Cambridge, MA, 2005).
- [41] S. Daulton, M. Balandat, and E. Bakshy, Parallel Bayesian optimization of multiple noisy objectives with expected hypervolume improvement, in *Advances in Neural Information Processing Systems*, edited by M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan (Curran Associates, Inc., Red Hook, NY, 2021), Vol. 34, pp. 2187–2200.

- [42] A. Ferran Pousa, S. Jalas, M. Kirchen, A. Martinez de la Ossa, M. Thévenet, S. Hudson, J. Larson, A. Huebl, J.-L. Vay, and R. Lehe, Bayesian optimization of laser-plasma accelerators assisted by reduced physical models, *Phys. Rev. Accel. Beams* **26**, 084601 (2023).
- [43] S. Hudson, J. Larson, J.-L. Navarro, and S. M. Wild, libEnsemble: A library to coordinate the concurrent evaluation of dynamic ensembles of calculations, *IEEE Trans. Parallel Distrib. Syst.* **33**, 977 (2022).
- [44] S. Schulz, I. Grguraš, C. Behrens, H. Bromberger, J. T. Costello, M. K. Czwalinna, M. Felber, M. C. Hoffmann, M. Ilchen, H. Y. Liu, T. Mazza, M. Meyer, S. Pfeiffer, P. Predki, S. Schefer, C. Schmidt, U. Wegner, H. Schlarb, and A. L. Cavalieri, Femtosecond all-optical synchronization of an x-ray free-electron laser, *Nat. Commun.* **6**, 5938 (2015).
- [45] R. Shalloo, C. Arran, G. Cheung, L. Corner, J. Holloway, R. Walczak, S. Hooker, N. Booth, O. Chekhlov, C. Gregory *et al.*, Measurement of femtosecond-scale drift and jitter of the delay between the North and South beams of Gemini, Central Laser Facility Annual Report 2015–2016, p. 32, https://www.clf.stfc.ac.uk/Pages/CLF_2015_16_Final.pdf.
- [46] J. Christie. Femtosecond synchronisation for externally-injected laser wakefield acceleration at CLARA, Ph.D. thesis, University of Liverpool, 2024.
- [47] P. Winkler, M. Trunk, L. Hübner, A. M. d. l. Ossa, S. Jalas, M. Kirchen, I. Agapov, S. A. Antipov, R. Brinkmann, T. Eichner, A. F. Pousa, T. Hülsenbusch, G. Palmer, M. Schnepf, K. Schubert, M. Thévenet, P. A. Walker, C. Werle, W. P. Leemans, and A. R. Maier, Active energy compression of a laser-plasma electron beam, *Nature (London)* **640**, 907 (2025).
- [48] A. Amodio, D. Wang, C. Berger, H.-E. Tsai, S. K. Barber, J. van Tilborg, A. Picksley, Z. Eisentraut, N. R. Vora, M. Logantha *et al.*, Pointing stabilization of a 1 Hz high-power laser via machine learning, *High Power Laser Sci. Eng.* **13**, e35 (2025).
- [49] M. Ma, M. Zeng, J. Wang, G. Lu, W. Yan, L. Chen, and D. Li, Particle-in-cell simulation of laser wakefield accelerators with oblique lasers in quasicylindrical geometry, *Phys. Rev. Accel. Beams* **28**, 021301 (2025).
- [50] S. Di Mitri, On the importance of electron beam brightness in high gain free electron lasers, *Photonics* **2**, 317 (2015).
- [51] Z. Huang and K.-J. Kim, Review of x-ray free-electron laser theory, *Phys. Rev. ST Accel. Beams* **10**, 034801 (2007).
- [52] K. Tiedtke *et al.*, The soft x-ray free-electron laser FLASH at DESY: Beamlines, diagnostics and end-stations, *New J. Phys.* **11**, 023029 (2009).
- [53] N. Bourgeois, J. Cowley, and S. M. Hooker, Two-pulse ionization injection into quasilinear laser wakefields, *Phys. Rev. Lett.* **111**, 155004 (2013).
- [54] L.-L. Yu, E. Esarey, C. B. Schroeder, J.-L. Vay, C. Benedetti, C. G. R. Geddes, M. Chen, and W. P. Leemans, Two-color laser-ionization injection, *Phys. Rev. Lett.* **112**, 125001 (2014).
- [55] S. M. Hooker, R. Bartolini, S. P. D. Mangles, A. Tünnemann, L. Corner, J. Limpert, A. Seryi, and R. Walczak, Multi-pulse laser wakefield acceleration: a new route to efficient, high-repetition-rate plasma accelerators and high flux radiation sources, *J. Phys. B* **47**, 234003 (2014).
- [56] O. Jakobsson, S. M. Hooker, and R. Walczak, GeV-scale accelerators driven by plasma-modulated pulses from kilohertz lasers, *Phys. Rev. Lett.* **127**, 184801 (2021).
- [57] J. Chappell, Simulation input scripts for “High brightness, symmetric electron bunch generation in a plasma wakefield accelerator via a radially-polarized plasma photocathode” [dataset], Zenodo, 2025, [10.5281/zenodo.15464164](https://doi.org/10.5281/zenodo.15464164).
- [58] S. Diederichs, C. Benedetti, A. Huebl, R. Lehe, A. Myers, A. Sinn, J.-L. Vay, W. Zhang, and M. Thévenet, HiPACE++: A portable, 3D quasi-static particle-in-cell code, *Comput. Phys. Commun.* **278**, 108421 (2022).