

Tuning Curves for a Laser-Plasma Accelerator

S. J alas,^{1,*} M. Kirchen,¹ C. Braun,¹ T. Eichner,¹ J. B. Gonzalez,¹ L. Hübner,¹ T. Hülsenbusch,¹ P. Messner,¹ G. Palmer,¹ M. Schnepf,² C. Werle,¹ P. Winkler,¹ W. P. Leemans,^{1,2} and A. R. Maier¹

¹*Deutsches Elektronen Synchrotron DESY, Notkestr. 85, 22607 Hamburg, Germany*

²*Department of Physics Universität Hamburg, Luruper Chaussee 149, 22761 Hamburg, Germany*

Applications of laser-plasma accelerators (LPA) require independent control of electron beam parameters. However, due to the complex coupling of the many variables governing the laser-plasma interaction, precisely tuning these parameters based on simple scalings is often impossible or at least sub-optimal. Here, we apply multi-objective Bayesian optimization to derive optimal tuning curves for LPAs, both in simulations and experiments. For electron energies between 150 and 250 MeV, we demonstrate tuning of the charge over a range of nearly 100 pC, while preserving optimal beam loading conditions with energy spreads below 5%. The derived tuning curves can explain the sometimes counter-intuitive interplay between laser and plasma control variables that is necessary to find the best trade-off between competing beam properties.

Over the past few years, laser-plasma accelerators (LPA) [1, 2] have shown rapid progress [3–5] promising compact and cost-effective drivers for medical [6, 7], high-energy physics [8, 9], and photon science [10–12] applications. Past experiments often focused on showcasing the enormous potential of LPAs with record beam properties. Real-world applications, however, will in addition require to precisely and independently control beam parameters over an extended tuning range. LPAs have previously demonstrated some tunability [13] but reaching the desired level of control is a very challenging task: Multiple tuning parameters can act directly or indirectly on the same beam property, and beam parameters may influence or even counteract each other through the complex laser-plasma interaction. For example, beam loading [14–17] couples bunch charge and energy spread. Simply varying one property (e.g., higher charge) can easily deteriorate (e.g., higher energy spread) the other. Consequently, tuning the electron beam phase space represents a multi-objective optimization problem, with the goal to identify solutions that provide the best trade-offs between multiple, conflicting objectives. The set of all of these optimal trade-offs forms a hypersurface through the space of possible solutions that is commonly referred to as the Pareto front. Gaining insight into the Pareto front is crucial for achieving optimal beam properties, as it helps to identify and reconcile the underlying conflicting mechanisms.

In this Letter, we present Pareto-optimal tuning curves to control the energy, charge, and energy spread in our laser-plasma accelerator. We utilize multi-objective Bayesian optimization (MOBO) [18–20] to sample the system’s Pareto front and derive tuning curves that provide accurate control over the relevant beam properties. We demonstrate the concept using simulations and present a method to handle the added complexity from noise under experimental conditions. Our experiments show that we can maintain optimal beam loading conditions over a broad range of beam configurations, allowing for flexible charge adjustments while minimizing negative impacts on energy spread at different specific design en-

ergies.

Bayesian optimization [21] (BO) is a method for efficiently finding global extrema in black box functions that are costly to evaluate and potentially noisy. Recently, it was explored as a way to find optimal operation points in conventional [22–24] and plasma accelerators [25, 26]. To avoid costly evaluations of the black box function, e.g., the accelerator, BO builds a surrogate model (typically a Gaussian process [27] model) of the system that allows to make cheap and noise-free predictions. With the surrogate model, an acquisition function is defined, which encodes the search strategy to identify the next point for physical evaluation. The acquired (experimental) data is used to refine the surrogate model and further guide the optimization.

In multi-objective optimization problems, there is not a unique optimal solution, but the goal is to find optimal trade-offs between many objectives, which define the Pareto front. A trade-off is considered optimal, or non-dominated, if it cannot be improved in one objective without deteriorating another. To find a set of solutions that is approaching the system’s Pareto front, MOBO maximizes the hypervolume that is spanned by the non-dominated solutions. For this, MOBO uses the expected improvement of the hypervolume as the acquisition function. As the hypervolume increases, the non-dominated solutions converge to the true Pareto front of the system.

So far, application of MOBO has been limited to simulation studies that recently have demonstrated its feasibility for optimizing conventional [28] and laser-plasma [29] accelerator setups. Implementation of the method under experimental conditions adds significant complexity due to measurement noise and shot-to-shot variations of the laser, and has not yet been shown successfully.

In the following, we derive a method to construct tuning curves for our laser-plasma accelerator based on MOBO. First, we illustrate our approach using particle-in-cell simulations [30] with the code FBPIC [31, 32], based on the LUX laser-plasma accelerator operated at DESY. Our setup is schematically shown in Fig. 1. The plasma

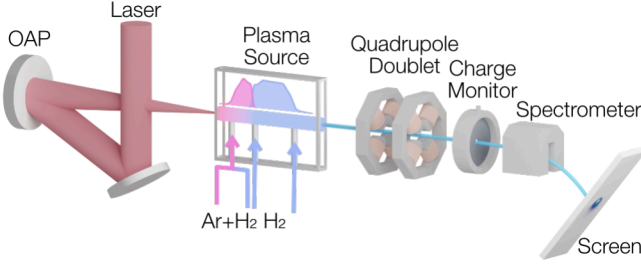


FIG. 1. LPA Setup: The drive laser has a variable energy and focus position. The plasma source has a mixed gas (magenta) and a pure Hydrogen (blue) filled region, which is controlled by four mass flow controllers. The accelerated beams are focused into a set of diagnostics, to measure beam charge and energy spectrum.

interaction is driven by the Ti:sapphire laser system ANGUS. The laser pulses (2.5 J, 36 fs FWHM) are focused to a spot size of $24\text{ }\mu\text{m}$ FWHM by a $f/25$ off-axis parabolic mirror (OAP). The laser energy is controlled by a motorized waveplate and a polarizer. The longitudinal focus position is fine-tuned by driving a lens in a telescope at the end of the laser chain. The plasma source, similar to the one used in [17, 26], is a capillary-like structure with three gas inlets that allow to localize a mixture of Argon and Hydrogen for electron injection in the front section, while providing a plateau of pure Hydrogen for acceleration in the back of the gas profile. The flow through the inlets is controlled with four individual mass flow controllers, two for the mixed gas inlet (Ar, H_2) and one for each of the inlets that make up the plateau (H_2). After the plasma exit, the electron beams are focused into a spectrometer (energy resolution: 0.1 %) with a pair of quadrupole magnets. A cavity-based charge monitor measures beam charge.

Starting from a working point of optimal beam loading conditions with minimum energy spread [17, 26], we want to tune the beam charge while preserving low energy spread. A simple way to vary the bunch charge in our setup is to shift the focal plane of the drive laser: Typically, the focus is set to the second half of the plasma profile to maximize the wakefield strength in the acceleration region. Shifting the focus upstream towards the mixed gas region increases the injection volume (ionization of inner shell Ar electrons) and thus beam charge. This effect is shown in Fig. 2a, continuously shifting the focus upstream over a range of $300\text{ }\mu\text{m}$. Starting off with a beam that is optimally beam loaded with 50 pC at 0.7% energy spread, the charge increases as intended. However, with this simple increase in charge, the previously flattened wakefield becomes overloaded and deforms, causing the bunch tail to experience a lower accelerating gradient than the head. As a result, the energy spread rapidly grows to unacceptable levels.

In principle, it is possible to counter-act this effect

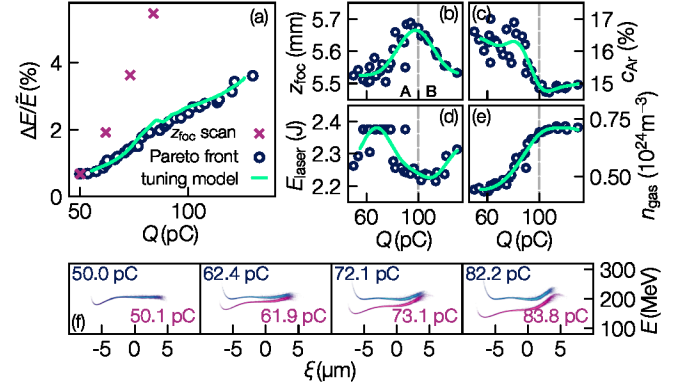


FIG. 2. LPA Tuning Curves (Simulations): The LPA setup is optimized for both charge and energy spread at $200 \pm 10\text{ MeV}$ using MOBO. Panel (a): energy spread, tuning the charge with the laser focus position only (crosses), or by a combined complex tuning of many parameters (circles), which are individually shown in panels (b)-(e). (A) and (B) mark two different tuning strategies. Panel (f): longitudinal phase spaces for the optimized tuning (blue) and the focus scan (violet). Solid lines, panels (a)-(e) mark the continuous tuning curves, derived from the optimization.

by increasing the wakefield strength. However, parameters that control the wakefield do not act in an isolated manner: For instance, raising the plasma density in return affects self-focusing and the evolution of the drive laser. Further adjustments of the laser energy and focus position would be required to recover the injection conditions; yet, they also influence the shape of the accelerating wakefield, which in turn could make a different bunch current profile necessary to achieve optimal beam loading. This requires adjustments of the Ar-concentration, which not only changes the injection rate, but also the current profile by affecting the plasma profile in the mixed gas region [26].

This cascade of consequences illustrates, that manually searching the functional dependencies that result in tuning of the beam charge while maintaining optimal beam loading and energy spread is futile. Instead, we apply MOBO to systematically study the input parameters that make up this optimal tuning behavior. Here, we use the tool BOTORCH [19, 33], and chose beam charge Q , relative energy spread $\Delta E/E$ (median absolute deviation), and the deviation from a design energy of 200 MeV, as objectives for optimization.

The results of our optimization, consisting of 700 individual simulations, are shown in Fig. 2a. Compared to the simple focus scan, we can increase the charge while keeping the energy spread low. As expected, to achieve this Pareto-optimal tuning, the input parameters (laser energy, focus position, gas density, and Ar concentration) need to be varied in non-obvious ways as shown in Fig. 2b-e.

Notably, for lower charges (region A) the focus posi-

tion z_{foc} and Ar concentration c_{Ar} have to be tuned in opposite directions from what would be expected when increasing the charge with these parameters individually. Instead, the focus is moved away from the mixed gas region, the Ar concentration is reduced, and gas density and laser energy are increased. The higher gas density increases the plasma density and the amount of Ar. Together with the higher laser energy, this results in an increase of the injected charge and a modification of the wakefield strength, which prevents excessive beam loading. Shifting the focus downstream prevents diffraction of the drive laser towards the end of the plasma and maintains optimal beam loading throughout the acceleration process. Moreover, a slight decrease in Ar concentration is necessary to match the injected current profile to the wakefield at the increased plasma density.

The tuning behavior changes completely once the gas density approaches the upper limit of the tuning range ($0.72 \times 10^{24} \text{cm}^{-3}$) at around 100 pC (region B). Now, to even further increase the charge, the focus needs to be tuned in the opposite direction, i.e., upstream, and the Argon concentration needs to be increased again slightly. The laser energy that was previously also reaching the maximum of the permitted range is slightly reduced.

As shown in Fig. 2f, this way of precisely balancing the input parameters maintains a flattened phase space over a wide range of charges, while keeping the beam energy fixed. The residual energy spread increase is primarily caused by correlated energy spread, building up only in the head and tail of the bunch. In contrast, the beams from the simple focus scan develop a significant energy correlation and energy deviation.

To derive an actual tuning curve that can be used in practice, the discrete, optimized working points have to be generalized. We fitted a multivariate kernel ridge regression model [34] to the data, that maps the desired beam charge Q to the set of input parameters \mathbf{X} , i.e., laser energy, focus position, gas density, and Ar concentration, that provide this charge at Pareto-optimal energy spread: $\mathbf{X}(Q)$. The model is shown in Fig. 2b-e, with the solid line indicating the continuous variation in input parameters as we continuously increase the beam charge. Using these tuning curves, the beam energy spread closely follows the Pareto front, compare Fig. 2a (solid line). Following our tuning curve, we can precisely set the beam charge, while providing the best possible energy spread and retaining the design energy.

Realizing the same concept experimentally is significantly more challenging, due to added measurement noise and shot-to-shot variations (jitter) in the experimental conditions.

Conceptually, we need to distinguish several sources of noise. First, the optimizer might set the machine to a certain target state (e.g, laser energy and focus position), that could be different for the actual shot. We can account for this with online shot-to-shot measurements

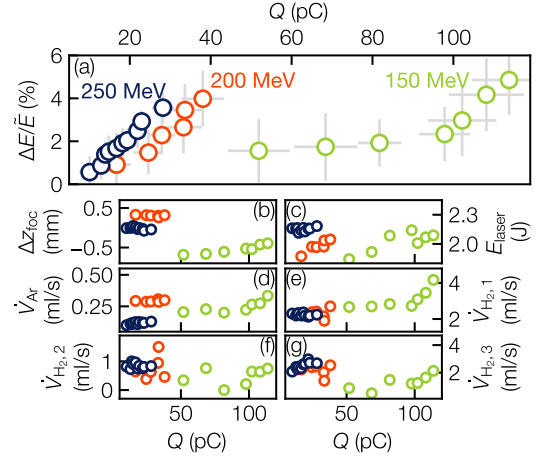


FIG. 3. Experimental results: (a) Pareto fronts of the beam charge and energy spread for beams within $\pm 10 \text{MeV}$ of the target energies with corresponding input parameters (b-g), i.e. the shift of the focus position Δz_{foc} , laser energy E_{laser} and the flows of the gas supplies \dot{V} .

of the laser parameters. Second, all measurements suffer from at least some measurement noise, that we cannot avoid. This random noise leads to uncertainty in the predictions of the surrogate model that can only be mitigated with additional data.

To cope with these effects, we recorded 40 shots at each setpoint. If we would simply take the average of the electron beam data from these shots, we would be averaging over the shot-to-shot variations of the laser and effectively treat them as random noise, which would lead to additional uncertainty of the surrogate model. Instead, at each working point, we trained a local Gaussian process model using the online measurements of the laser energy and focus position as inputs and the electron parameters as outputs. This allowed us to disentangle the effects of the laser jitter and interpolate the properties of the electron beam for the exact input parameters that were requested by the optimizer. Then we used this information to train the surrogate model that guided the optimization. While previous experiments [26] trained the surrogate model directly on all individual shots, the computational cost associated with the MOBO algorithm when dealing with large datasets led us to adopt this approach of local modeling.

To explore the capabilities of our setup, we performed a series of experiments with MOBO runs at 150, 200 and 250 MeV. As before, the objectives for the optimization were the beam charge, relative energy spread, and deviation from the reference energy. The resulting Pareto fronts with their corresponding input parameters after 100 iterations, i.e., 4000 individual shots, are shown in Fig. 3.

For all three target energies we find low energy spread setpoints over a significant range of beam charges. Con-

sistent with the simulations, the energy spread grows with increasing charge. For lower target energies, the achievable charge increases as one would expect from an aspect of energy conservation. As a result, for 150 MeV, it was possible to tune the beam charge between 50 and 130 pC, while maintaining a relative energy spread of less than 5%. Up to a charge of 100 pC, the energy spread only shows a modest growth and stays below 2.5% but rapidly increases from there on. This behavior implies that the available input parameter range did not support optimal beam loading beyond this point. Therefore, additional charge can only be loaded into the wakefield at the cost of significantly distorting the accelerating field. For the two higher energy cases, this quick increase of the energy spread is present already at low charge, indicating that here beam loading conditions immediately worsen so much with increasing beam charge that a significant energy correlation is imprinted.

Fig. 3b-g shows, that the three different energies populate very different islands in the input parameter space, especially for Δz_{foc} , \dot{V}_{Ar} and $\dot{V}_{\text{H}_{2,3}}$. Also the way some parameters are tuned varies from energy to energy. For example, for 150 MeV, the focus is moved downstream to increase the charge similar to the lower charged beams in the simulated setup before. For the other two cases it is moved upstream towards the mixed gas comparable to the higher charge cases in the simulations. This again hints that these two cases are already in a regime that is dominated by strong beam loading. Assuming this is true, we expect that these cases would behave similar to the 150 MeV case when provided with more laser energy and exhibit lower energy spread over a wider tuning range. The fact that these three cases behave so diverse in regard to the input parameters highlights the complexity of, first, finding beams with acceptable quality and then, second, being able to control them.

A general feature that is shared by nearly all configurations in Fig. 3 is that the flow through the last gas inlet $\dot{V}_{\text{H}_{2,3}}$ is higher than through the middle one $\dot{V}_{\text{H}_{2,2}}$. By creating a slight density ramp the slippage between the electrons and the laser can be compensated through contraction of the wakefield which is a way to increase efficiency and load more charge.

The Pareto fronts in our experiment, Fig. 3, have fewer sample points compared to the previous simulations, as a result of the optimizer being run for a lower number of iterations. Consequently, providing a continuous tuning curve is even more important. The data shows a less non-linear behavior than the simulations before, possibly because it is less converged to the true Pareto front of the system. In this case, it is sufficient to use a simple multivariate linear model to describe the main parameter variations. We derived such a model, exemplary for

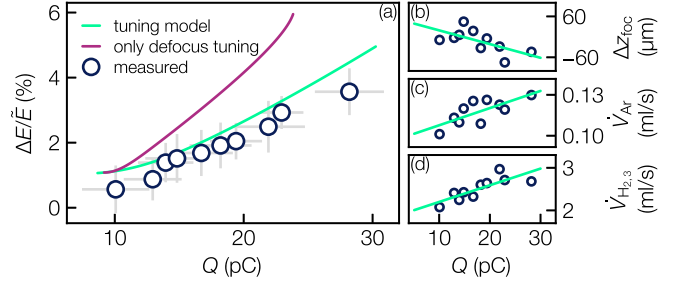


FIG. 4. Experiment: Beam charge tuning curve at 250 MeV. A linear model (green line) is fitted to the optimization result mapping the requested beam charge to the corresponding input parameters along the Pareto front (b-d). (a) The tuning curve was validated and compared to the measured data (circles) with predictions from a Gaussian process surrogate model. For comparison the effect of a laser focus position scan over 450 μm is shown (red line).

250 MeV,

$$\begin{aligned} \mathbf{X}(Q) &= (\Delta z_{\text{foc}}(Q), \dot{V}_{\text{Ar}}(Q), \dot{V}_{\text{H}_{2,3}}(Q)) \\ &\approx ([59 - 4Q] \mu\text{m}, [\frac{9.5}{10^2} + \frac{1.25}{10^3} Q] \frac{\text{ml}}{\text{s}}, [1.8 + \frac{4}{10^2} Q] \frac{\text{ml}}{\text{s}}), \end{aligned}$$

with Q in units of pC. Here we only considered parameters that made significant contributions to the tuning of the beam charge, keeping the remaining inputs constant ($E_{\text{laser}} = 2.15 \text{ J}$, $\dot{V}_{\text{H}_{2,1}} = 2.24 \text{ ml/s}$, $\dot{V}_{\text{H}_{2,2}} = 0.84 \text{ ml/s}$). For example, to generate beams with 10 pC, according to the tuning curve, an operator would set the machine to $\Delta z_{\text{foc}} = 19 \mu\text{m}$, $\dot{V}_{\text{Ar}} = 0.107 \text{ ml/s}$ and $\dot{V}_{\text{H}_{2,3}} = 2.2 \text{ ml/s}$.

The input parameters over the full range of the tuning curve and the underlying data is shown in Fig. 4b-d. The tuning strategy is to shift the focus towards the mixed gas region and to increase the flow of Ar to enhance the beam charge. To support the additional charge and balance beam loading, the flow of H_2 in the back of the plasma source needs to go up. This leads to a higher gas pressure in the plateau region and creates a density ramp towards the back of the plasma profile, ensuring optimal performance.

To evaluate the effectiveness of the tuning curve, we employ the surrogate model obtained from the optimization process to forecast the response of the LPA when adjusted based on the curve. For the example above, the surrogate model predicts that the LPA generates beams with $(10.9 \pm 1.5) \text{ pC}$ at $(1.2 \pm 0.4) \%$ energy spread that are within $(4.3 \pm 1.8) \text{ MeV}$ of the 250 MeV design energy.

The green line in Fig. 4a shows the predicted behavior of the energy spread when scanning the charge according to the tuning curve. The resulting working points closely align with those measured during optimization, indicating their proximity to the presumed Pareto front of the system. Remarkably, this behavior is achieved using only the three most important input parameters

identified through our analysis. In contrast, when tuning the charge only with the focus position, the resulting growth of energy spread is roughly a factor of two larger than what is obtained with our tuning curve. Provided with the simple tuning curve, operators can precisely set the beam charge required for their application, while still achieving the smallest possible energy spread.

In conclusion, we have demonstrated dedicated control over electron bunch parameters in our laser-plasma accelerator using continuous tuning curves, both in simulations and experiments, that have been build via multi-objective Bayesian optimization of the machine. In particular, the bunch charge was varied over a range of nearly 100 pC and at different electron energies (150-250 MeV), while providing the smallest possible energy spreads ($<5\%$) by moving the laser and plasma control parameters along Pareto-optimal tuning curves. Even more importantly, we could show that the derived tuning curves itself are meaningful expressions of the control and output parameter relationships, which allowed us to identify and explain how different physics mechanisms have to be balanced to provide optimal trade-offs between competing beam parameters.

While we showcase the potential of our our method with our own LPA, our approach is generally valid and can be readily applied to other laser-plasma accelerator setups: We have introduced MOBO as a powerful tool to study the complex interplay of the many mechanisms that eventually result in a high-quality laser-plasma electron beams. We believe, this method will be broadly adopted in the future.

To drive applications, laser-plasma accelerators need to be operated with tunability over a broad range of parameters, while still providing sub-percent-level energy spread beams. In our experiment, we were limited to sub-5% energy spread beams. However, studying the carefully balanced mechanism at the Pareto front indicates, that with additional control parameters, and a larger range of laser energy and plasma density, we should be able to extend the range of sub-percent energy spread beams. We expect that in the future, MOBO will be used as an important tool to study and design laser-plasma accelerators with the degrees of freedom (control parameters) required to achieve extended tuning ranges. Finally, tuning curves, as we present here, will allow operators to precisely and quickly configure the electron beam properties, which will be crucial for applications.

We appreciate the support by the workshops and technical groups at DESY. This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 491245950.

* soeren.jalas@desy.de

- [1] T. Tajima and J. M. Dawson, “Laser electron accelerator,” *Phys. Rev. Lett.* **43**, 267 (1979).
- [2] E. Esarey, C. B. Schroeder, and W. P. Leemans, “Physics of laser-driven plasma-based electron accelerators,” *Rev. Mod. Phys.* **81**, 1229 (2009).
- [3] A. J. Gonsalves, K. Nakamura, J. Daniels, C. Benedetti, C. Pieronek, T. C. H. de Raadt, S. Steinke, J. H. Bin, S. S. Bulanov, J. van Tilborg, C. G. R. Geddes, C. B. Schroeder, Cs. Tóth, E. Esarey, K. Swanson, L. Fan-Chiang, G. Bagdasarov, N. Bobrova, V. Gasilov, G. Korn, P. Sasorov, and W. P. Leemans, “Petawatt laser guiding and electron beam acceleration to 8 gev in a laser-heated capillary discharge waveguide,” *Phys. Rev. Lett.* **122**, 084801 (2019).
- [4] A. R. Maier, N. M. Delbos, T. Eichner, L. Hübner, S. Jalas, L. Jappe, S. W. Jolly, M. Kirchen, V. Leroux, P. Messner, M. Schnepf, M. Trunk, P. A. Walker, C. Werle, and P. Winkler, “Decoding sources of energy variability in a laser-plasma accelerator,” *Phys. Rev. X* **10**, 031039 (2020).
- [5] L. T. Ke, K. Feng, W. T. Wang, Z. Y. Qin, C. H. Yu, Y. Wu, Y. Chen, R. Qi, Z. J. Zhang, Y. Xu, X. J. Yang, Y. X. Leng, J. S. Liu, R. X. Li, and Z. Z. Xu, “Near-gev electron beams at a few per-mille level from a laser wakefield accelerator via density-tailored plasma,” *Phys. Rev. Lett.* **126**, 214801 (2021).
- [6] L. Labate, D. Palla, D. Panetta, F. Avella, F. Baffigi, F. Brandi, F. Di Martino, L. Fulgentini, A. Giulietti, P. Köster, D. Terzani, P. Tomassini, C. Traino, and L. A. Gizzi, “Toward an effective use of laser-driven very high energy electrons for radiotherapy,” *Sci. Reports* **10**, 17307 (2020).
- [7] K. Svendsen, D. Guenot, J.B. Svensson, K. Petersson, A. Persson, and O. Lundh, “A focused very high energy electron beam for fractionated stereotactic radiotherapy,” *Sci. Reports* **11**, 5844 (2021).
- [8] C. B. Schroeder, E. Esarey, C. G. R. Geddes, C. Benedetti, and W. P. Leemans, “Physics considerations for laser-plasma linear colliders,” *Phys. Rev. ST Accel. Beams* **13**, 101301 (2010).
- [9] C. Benedetti, S. S. Bulanov, E. Esarey, C. G. R. Geddes, A. J. Gonsalves, A. Huebl, R. Lehe, K. Nakamura, C. B. Schroeder, D. Terzani, J. van Tilborg, M. Turner, J. L. Vay, T. Zhou, F. Albert, J. Bromage, E. M. Campbell, D. H. Froula, J. P. Palastro, J. Zuegel, D. Bruhwiler, N. M. Cook, B. Cros, M. C. Downer, M. Fuchs, B. A. Shadwick, S. J. Gessner, M. J. Hogan, S. M. Hooker, C. Jing, K. Krushelnick, A. G. R. Thomas, W. P. Leemans, A. R. Maier, J. Osterhoff, K. Poder, M. Thevenet, C. Joshi, W. B. Mori, H. M. Milchberg, M. Palmer, J. G. Power, and N. Vafaei-Najafabadi, “Linear colliders based on laser-plasma accelerators,” (2022).
- [10] H. Schwöerer, B. Liesfeld, H.-P. Schlenvoigt, K.-U. Amthor, and R. Sauerbrey, “Thomson-backscattered x rays from laser-accelerated electrons,” *Phys. Rev. Lett.* **96**, 014802 (2006).
- [11] S. Kneip, C. McGuffey, J. L. Martins, S. F. Martins, C. Bellei, V. Chvykov, F. Dollar, R. Fonseca, C. Huntington, G. Kalintchenko, A. Maksimchuk, S. P. D. Mangles, T. Matsuoka, S. R. Nagel, C. A. J. Palmer, J. Schreiber, K. Ta Phuoc, A. G. R. Thomas, V. Yanovsky, L. O. Silva, K. Krushelnick, and Z. Najmudin, “Bright spatially coherent synchrotron x-rays from a table-top source,” *Nature Physics* **6**, 980–983 (2010).

- [12] W. Wang, K. Feng, L. Ke, C. Yu, Y. Xu, R. Qi, Y. Chen, Z. Qin, Z. Zhang, M. Fang, J. Liu, K. Jiang, H. Wang, C. Wang, X. Yang, F. Wu, Y. Leng, J. Liu, R. Li, and Z. Xu, “Free-electron lasing at 27 nanometres based on a laser wakefield accelerator,” *Nature* **595**, 516–520 (2021).
- [13] A. J. Gonsalves, K. Nakamura, C. Lin, D. Panasenkov, S. Shiraishi, T. Sokollik, C. Benedetti, C. B. Schroeder, C. G. R. Geddes, J. van Tilborg, J. Osterhoff, E. Esarey, C. Toth, and W. P. Leemans, “Tunable laser plasma accelerator based on longitudinal density tailoring,” *Nature Physics* **7**, 862–866 (2011).
- [14] 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).
- [15] J. P. Couperus, R. Pausch, A. Köhler, O. Zarini, J. M. Krämer, M. Garten, A. Huebl, R. Gebhardt, U. Helbig, S. Bock, K. Zeil, A. Debus, M. Bussmann, U. Schramm, and A. Irman, “Demonstration of a beam loaded nanocoulomb-class laser wakefield accelerator,” *Nat. Commun.* **8**, 487 (2017).
- [16] J. Götzfried, A. Döpp, M. F. Gilljohann, F. M. Foerster, H. Ding, S. Schindler, G. Schilling, A. Buck, L. Veisz, and S. Karsch, “Physics of high-charge electron beams in laser-plasma wakefields,” *Phys. Rev. X* **10**, 041015 (2020).
- [17] M. Kirchen, S. Jalas, P. Messner, P. Winkler, T. Eichner, L. Hübner, T. Hülsenbusch, L. Jeppe, T. Parikh, M. Schnepf, and A. R. Maier, “Optimal beam loading in a laser-plasma accelerator,” *Phys. Rev. Lett.* **126**, 174801 (2021).
- [18] M.T.M. Emmerich, K.C. Giannakoglou, and B. Naujoks, “Single- and multiobjective evolutionary optimization assisted by gaussian random field metamodels,” *IEEE Transactions on Evolutionary Computation* **10**, 421–439 (2006).
- [19] 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* (2020).
- [20] 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*, Vol. 34 (2021) pp. 2187–2200.
- [21] D. R. Jones, M. Schonlau, and W. J. Welch, “Efficient global optimization of expensive black-box functions,” *J. Global Optim.* **13**, 455 (1998).
- [22] J. Duris, D. Kennedy, A. Hanuka, J. Shtalenkova, A. Edelen, P. Baxevanis, A. Egger, T. Cope, M. McIntire, S. Ermon, and D. Ratner, “Bayesian optimization of a free-electron laser,” *Phys. Rev. Lett.* **124**, 124801 (2020).
- [23] J. Kirschner, M. Mutný, A. Krause, J. Coello de Portugal, N. Hiller, and J. Snurverink, “Tuning particle accelerators with safety constraints using bayesian optimization,” *Phys. Rev. Accel. Beams* **25**, 062802 (2022).
- [24] C. Xu, T. Boltz, A. Mochihashi, A. Santamaria Garcia, M. Schuh, and A.-S. Müller, “Bayesian optimization of the beam injection process into a storage ring,” *Phys. Rev. Accel. Beams* **26**, 034601 (2023).
- [25] R. J. Shalloo, S. J. D. Dann, J.-N. Gruse, C. I. D. Underwood, A. F. Antoine, C. Arran, M. Backhouse, C. D. Baird, M. D. Balcazar, N. Bourgeois, J. A. Cardarelli, P. Hatfield, J. Kang, K. Krushelnick, S. P. D. Mangles, C. D. Murphy, N. Lu, J. Osterhoff, K. Pöder, P. P. Rajeev, C. P. Ridgers, S. Rozario, M. P. Selwood, A. J. Shahani, D. R. Symes, A. G. R. Thomas, C. Thornton, Z. Najmudin, and M. J. V. Streeter, “Automation and control of laser wakefield accelerators using bayesian optimization,” *Nat. Commun.* **11**, 6355 (2020).
- [26] S. Jalas, M. Kirchen, P. Messner, P. Winkler, L. Hübner, J. Dirkwinkel, M. Schnepf, R. Lehe, and A. R. Maier, “Bayesian optimization of a laser-plasma accelerator,” *Phys. Rev. Lett.* **126**, 104801 (2021).
- [27] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning*. (MIT Press, 2006).
- [28] R. Roussel, A. Hanuka, and A. Edelen, “Multiobjective bayesian optimization for online accelerator tuning,” *Phys. Rev. Accel. Beams* **24**, 062801 (2021).
- [29] F. Irshad, S. Karsch, and A. Döpp, “Multi-objective and multi-fidelity bayesian optimization of laser-plasma acceleration,” *Phys. Rev. Res.* **5**, 013063 (2023).
- [30] The simulations had a grid resolution of $\Delta z = 0.02 \mu\text{m}$, $\Delta r = 0.3 \mu\text{m}$ with 8 particles per cell and used quasi-cylindrical geometry with two azimuthal modes.
- [31] 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).
- [32] M. Kirchen, R. Lehe, B. B. Godfrey, I. Dornmair, S. Jalas, K. Peters, J.-L. Vay, and A. R. Maier, “Stable discrete representation of relativistically drifting plasmas,” *Phys. Plasmas* **23**, 100704 (2016).
- [33] M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy, “Botorch: A framework for efficient monte-carlo bayesian optimization,” in *Advances in Neural Information Processing Systems* **33** (2020).
- [34] K. P. Murphy, *Machine learning: a probabilistic perspective* (MIT press, 2012) pp. 492–493.