

Data-driven modeling of a laser-plasma accelerator-based x-ray source

F. Brogren¹, S. J alas², L. Hübner², P. Messner², M. Schnepf³, M. Trunk², C. Werle²,
 P. Winkler², M. Marklund¹, A. Gonoskov¹, W. P. Leemans^{2,3},
 A. R. Maier² and M. Kirchen²

¹*Department of Physics, University of Gothenburg, Origovägen 6B, 412 96 Gothenburg, Sweden*

²*Deutsches Elektronen-Synchrotron DESY, Notkestraße 85, 22607 Hamburg, Germany*

³*Department of Physics, Universität Hamburg, Luruper Chaussee 149, 22761 Hamburg, Germany*



(Received 28 March 2025; accepted 26 September 2025; published 21 October 2025)

Laser-plasma accelerators (LPAs) enable compact, bright x-ray sources, but their practical application demands a significant reduction of beam instabilities that originate from drive laser fluctuations. Moreover, the complexity of the laser-plasma interaction makes it difficult to disentangle and quantify the impact of individual parameters on machine performance. To address this challenge, we develop a data-driven modeling strategy and apply it to an extensive dataset collected during a daylong operation of the LPA-based x-ray source LUX. By making use of an orthogonal-distance-based training objective, our approach reduces bias originating from measurement errors, which allows us to estimate the functional dependencies between laser, accelerated electrons, and generated undulator radiation. In addition, we demonstrate accurate prediction of the x-ray spectrum based on noninvasive measurements, showcasing the potential of our approach for building virtual diagnostics.

DOI: [10.1103/dnnw-lfn5](https://doi.org/10.1103/dnnw-lfn5)

I. INTRODUCTION

Laser-plasma accelerators (LPA) [1,2] can deliver high-energy [3] and high-current [4] electron beams within only a few centimeters. These compact sources are therefore attractive for applications such as radiation therapy [5] and medical imaging [6], and, as shown in recent experiments, capable of driving free-electron lasers down to the vacuum ultraviolet regime [7–9]. Despite these milestones, widespread adoption of laser-plasma accelerators is still limited because they lag behind conventional technology in beam quality and shot-to-shot stability. This is primarily a consequence of the high sensitivity of the acceleration process to small variations of the drive laser pulse [10,11]. It is widely accepted that resolving this issue primarily depends on advancements in high-power laser technology.

However, the complexity of the nonlinear laser-plasma interaction makes it difficult to disentangle the influence of different laser parameters and identify those that have the largest impact on the electron beam properties. To enable targeted laser development and enhance the LPA performance, it is essential to develop a deeper understanding of the causal relationship between laser, electron, and secondary radiation parameters.

Carrying out this task with traditional data analysis has proven difficult due to inherent noise and measurement uncertainties in these experiments. Yet, recent progress in long-term stability and the push toward higher repetition rates have significantly improved the quality and availability of experimental data [12,13].

Given these developments, data-driven approaches have emerged as a powerful new tool for understanding and controlling the physics of plasma accelerators [14–16]. By leveraging machine learning algorithms, noisy experimental data can be converted into high-dimensional models that reveal the underlying system dynamics. Such models not only help in identifying the sources of electron variability but also provide the foundation for autonomous optimization or can serve as a virtual diagnostic [17]. Particularly for virtual diagnostics, interest is growing as they could enable applications where maintaining stability is less critical, as long as accurate information about the generated electron beam or secondary radiation is available.

In recent years, data-driven methods have become increasingly popular in the field of plasma acceleration. For instance, several experiments have demonstrated the use of Gaussian process modeling to improve the electron beam quality via Bayesian optimization [18,19]. Other studies have shown that it is possible to accurately predict the electron beam properties from measured laser data, using different model architectures, such as fully connected neural networks [11], variational neural networks [15], and random forests [20]. In addition, there are several reports of successful reconstruction of x-ray spectra using

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.

noninvasive measurements, although these studies were conducted at conventional accelerators [21–23].

So far, these data-driven methods have focused primarily on predictive accuracy rather than *explainability*, in regard to estimating the correlations of the underlying physical processes. Although prediction and estimation are used interchangeably in many fields, there is cause to differentiate between the two. Prediction generally refers to the “guessing” of the (future) response of some process, while estimation refers to the guessing of model parameters, both with the support of data [24,25].

To provide a concrete example, consider a process modeled by $y = f(x, \theta)$, where θ is a model parameter. In prediction tasks, the objective is to minimize the residuals between $f(x, \theta)$ and y by adjusting θ . Conversely, for parameter estimation, the objective shifts to minimizing the error of θ itself [26]. This is often motivated by some physical meaning of θ that describes the true dependence of y on x . In many cases, these two objectives align. However, the differentiation becomes important when there is a measurement error in the input variable x , which is a common scenario in experiments.

In this paper, we follow this line of thought and develop a versatile approach to data-driven modeling of experiments based on feedforward neural networks. Depending on the task, our framework allows targeting either one of the above objectives to derive an *explanatory model* or *predictive model*.

Both models are employed to gain specific insights into the physics of a LPA-driven x-ray source. The analysis is based on a large dataset generated during a daylong operation of the LUX laser-plasma accelerator [12], where high-quality electron beams were accelerated to 260 MeV and subsequently used to generate spontaneous undulator radiation centered at 12 nm wavelengths.

The paper is structured as follows. In Sec. II, we motivate our approach and present a detailed derivation of the modeling process. In Sec. III, we describe the experimental setup and present our feature selection process, which is quantified by *permutation importance* and low correlation between input variables. In Sec. IV, we demonstrate the usefulness of the explanatory model as a tool in finding the causal relation between input and output variables in the presence of measurement errors. Throughout several examples, we uncover the influence of individual laser properties on the LPA process and the generated undulator radiation. Finally, we shift our attention to the predictive model and demonstrate its effectiveness in building virtual diagnostics. By using only noninvasive measurements, we recreate the x-ray spectrum, achieving an R^2 score up to 0.98 for the central wavelength.

II. DATA-DRIVEN MODELING

As many experimental systems, the LPA process can be characterized by a mapping between measurable input

parameters and output quantities. Uncertainty in these variables arises from diagnostic inaccuracies and from unmeasured fluctuations in the system state, which effectively appear as noise on the observed output.

Using models that assume error-free inputs, e.g., ordinary least squares regression, on data with input errors generally leads to biased estimates, a phenomenon known as attenuation bias [27,28]. Several methods have been developed to address this issue, including Deming regression and its extensions [29,30], total least squares [31], instrumental variables [32], and repeated measurements [32,33]. However, most existing measurement error models rely on specific assumptions about the functional form, such as linearity in regression parameters. In contrast, models with more relaxed assumptions, e.g., spline-based methods, require the selection of multiple hyperparameters, which is increasingly problematic with higher dimensionality [34].

Feedforward neural networks (FNNs) can approximate complex, nonlinear functions in high-dimensional spaces, making them a natural choice for the type of data-driven inference considered here [35]. Although they make no explicit assumptions about the functional form, they still rely on implicit assumptions about the data-generating process, which can be critical for accurate modeling. This includes, in particular, assumptions about the error distribution.

To illustrate this point, we present a synthetic example in Fig. 1, where two neural networks are trained to recover a known underlying function (black) from noisy data. The model trained with mean squared error (MSE) (orange) fails to capture the correct relationship, since it is minimizing the




FIG. 1. Two models were trained to estimate the underlying dependence f_{true} (black) using different optimization criteria: orthogonal distance (OD) marked in blue and mean squared error (MSE) marked in orange. The randomly generated data points x and y had a normally distributed error with $\mu = 0$ and $\sigma = 0.7$. The black dots represent a few random samples of the validation set and the 2D-histogram shows the entire dataset. The shaded area represents the 95% quantile interval across the predictions from an ensemble of 50 estimators.

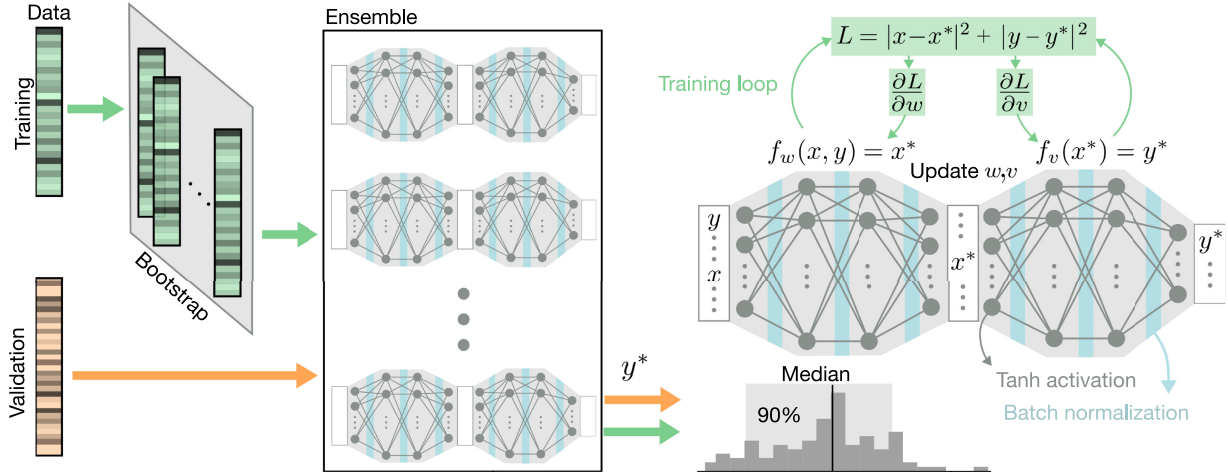


FIG. 2. Data is split into a training and a validation set. The training set is randomly sampled with replacement to create subsamples. Each subsample is given to one network, which is then trained by repeatedly propagating the data, evaluating the orthogonal distance and updating the weights through back propagation. The final model is calculated from the median of all networks in the ensemble and the model uncertainty from the quantiles of the distribution of models. Each network (upper right) consists of two FNNs. The first takes in measurement data x and y and generates an estimation of the x without measurement noise (x^*). The second FNN takes in x^* and generates an estimation of y without measurement noise (y^*). Once trained f_v will be an estimation of the underlying dependence between x and y .

vertical distance between model and data points. In cases where there are errors on both input (x) and output (y), the optimization criterion should instead minimize the orthogonal (or total) distance (OD) between model and data, as demonstrated by the improved reconstruction obtained with OD minimization (blue) in Fig. 1.

In the following Sec. II A, we present the implementation of such an OD optimization criteria in the context of neural networks. A similar approach as the one developed in this paper can be found in [34]. Note that the goal of this model is to estimate the dependencies between input and output of the system, not solely to predict the system output. This is an important distinction since it affects the optimization condition and consequently the modeling. In Sec. II B, we turn to the problem of prediction and elaborate on the distinction between prediction and estimation. We call the models resulting from these two approaches the *explanatory model* and *predictive model*.

A. Explanatory model

For formulating a mathematical description of the system, we define x_i and y_i as the vectors of measured input and output variables, where $i \in [1, N]$ and N are the number of measurements. Note that $x_i = [x_{i1}, \dots, x_{in_x}]^T$ and $y_i = [y_{i1}, \dots, y_{in_y}]^T$ are vectors where n_x, n_y are the number of input variables and output variables, respectively.

The problem of measurement errors can mathematically be expressed as

$$x_i = x_i^* + \varepsilon_{x,i}, \quad y_i = f(x_i^*) + \varepsilon_{y,i}, \quad (1)$$

where x_i^* are the true measurement values and $\varepsilon_{x,i}, \varepsilon_{y,i}$ are assumed to be normally distributed (multivariate) noise parameters with the same dimensionality as x_i, y_i . The goal is to estimate f . We model the system using two FNNs f_w, f_v with weights w, v , where $f_w(x_i, y_i)$ is an estimation of x_i^* and $f_v(f_w(x_i, y_i))$ is an estimation of $f(x_i^*)$. When optimized, the function f_v will be an estimation of f . A schematic figure of the model is shown in Fig. 2 (upper right). To find the values of v, w such that f_v, f_w align with the data, the likelihood is maximized, that is, maximizing the probability of obtaining the data $x = \{x_i\}, y = \{y_i\}$ given v, w : $P(x, y|v, w)$. The presence of measurement errors on output and input implies that for optimally trained networks the data $D_i = [x_i, y_i]$ should be distributed according to a multivariate Gaussian distribution centered around $\mu_{wv,i} = [f_w(x_i, y_i), f_v(f_w(x_i, y_i))]^T$. This assumption results in the likelihood

$$P(x, y|v, w) = \prod_i \frac{1}{\sqrt{(2\pi)^k \det(\Sigma)}} e^{-\frac{1}{2}[D_i - \mu_{wv,i}]^T \Sigma^{-1} [D_i - \mu_{wv,i}]}, \quad (2)$$

where Σ is the covariance matrix of the errors ε_x and ε_y with dimension $k \times k$, where $k = n_x + n_y$. Assuming that Σ is diagonal, that is, assuming no covariance between errors, the expression is simplified. Moreover, by taking the negative log of the expression and discarding all constant terms a more computer-friendly optimization condition for v, w is obtained

$$\min_{v,w} \sum_{i,j} \frac{(D_{ij} - \mu_{wv,ij})^2}{\sigma_j^2} = \min_{v,w} \sum_{i,j,j'} \frac{[y_{ij'} - f_v(f_w(x_i, y_i))]_j^2}{\sigma_{y,j'}^2} + \frac{[x_{ij} - f_w(x_i, y_i)]_j^2}{\sigma_{x,j}^2}, \quad (3)$$

where the index j and j' are running over all input and output variables, respectively, and $\sigma_{x,j}^2, \sigma_{y,j'}^2$ are the diagonal elements of Σ , that is the variance of the errors $\varepsilon_{x,j}, \varepsilon_{y,j'}$. Consequently, maximizing the likelihood is equivalent to minimizing the orthogonal distance between the model and the data. We refer to the term in Eq. (3), prior to minimization, as the OD loss. Note two things about the above expression: (i) The variances $\sigma_{x,j}^2, \sigma_{y,j'}^2$, which in the experiment corresponds to the variance of the measurement error, are free variables and need to be estimated. (ii) If x_i is measured with a small error compared to y_i , i.e., $\sigma_{x,j}^2 \ll \sigma_{y,j'}^2$ for all j, j' , it follows that $f_w(x_i, y_i)_j \approx x_{ij}$ since the last term in Eq. (3) would drastically increase otherwise. Consequently, the optimization criteria effectively transform into the MSE.

As a result, the OD loss includes the MSE optimization criteria. For MSE, however, assumptions regarding error distributions remain implicit, whereas they are explicitly stated in the OD loss.

Ideally, the measurement error for each input and output variable should be estimated individually and incorporated into the OD loss. However, obtaining accurate error estimates is often challenging in practice. In the context of experimental data, assuming equal measurement uncertainty for both input and output variables (i.e., setting $\sigma_{x,j} = \sigma_{y,j'}$ for all j, j') is arguably a more reasonable approximation than the one implicit in MSE minimization, which presumes $\sigma_{x,j}^2 \ll \sigma_{y,j'}^2$.

It is important to note that only the ratio between σ_x and σ_y affects the model outcome, not their absolute values. This follows directly from the structure of the OD loss Eq. (3), where uniform scaling of all terms leaves the position of the minimum unchanged, whereas changes in their relative weights do not.

Based on this reasoning, we assume that the variance of the measurement error is on the same scale as the root-mean-square deviation (RMSD) of the data. Since neural networks benefit from inputs and outputs being similarly scaled, the data are normalized to zero mean and unit RMSD prior to training. Under this normalization, the assumption of equal measurement errors is implemented by setting $\sigma_{x,j}^2 = \sigma_{y,j'}^2 = 1$ for all j, j' .

Further details on the choice of measurement error parameters are provided in Supplemental Material [36], where we also discuss the consequences of mismatched error assumptions. In particular, we show in an illustrative

example that the OD loss outperforms the MSE loss except in the limiting case where $\sigma_{x,j}^2 \ll \sigma_{y,j'}^2$.

Under optimal training, f_w is expected to approximate an inverse of f_v , in the sense that $f_w(x', f_v(x')) = x'$ where x' is an arbitrary vector. However, the underlying function (f) may not be bijective, leading to ambiguity when the gradient is close to zero, as for $x = 0$ in Fig. 1.

We found that this flexibility can result in sharp changes of f_w and large variation of function shapes between different training sessions. A common reason for high variance are large neural weight values; a widely utilized method to address this issue is weight regularization. In L2-regularization the weights are pushed toward lower values by introducing a penalizing term $\eta \sum_i |w_i|^2$ in the loss. In a Bayesian framework, it can be shown that L2-regularization is equivalent to imposing a Gaussian prior on the weights [37]. Note that η is a regularization parameter and must be chosen based on convergence and model performance. The training loop can be summarized as (i) One batch of data is propagated through the networks, which gives $f_w(x_i, y_i), f_v(f_w(x_i, y_i))$. (ii) The OD loss $L = \sum_{i,j} \frac{(D_{ij} - \mu_{wv,ij})^2}{\sigma_j^2} - \eta_w \sum_i |w_i|^2 - \eta_v \sum_i |v_i|^2$ is calculated. (iii) The partial derivative with respect to w, v of L is calculated. (iv) The network is updated by $w_+ = s_w \frac{\partial L}{\partial w} + w$, $v_+ = s_v \frac{\partial L}{\partial v} + v$, where s_w, s_v are the step lengths.

From a probabilistic view the training data is sampled from a larger distribution of data, therefore a model optimized on a subsample might differ from a model created on another subsample. This uncertainty is called model uncertainty or epistemic uncertainty and will decrease with the amount of data [38,39]. Model uncertainty can be approximated by creating subsamples of the dataset using random sampling (with replacement) and training models separately on these subsets, see Fig. 2. This technique is known as bootstrapping [40]. The variation of these models approximates the model uncertainty [41], furthermore this uncertainty will grow when the data is sparse, which makes the modeled uncertainty a good indication of the region where the model is trustworthy.

B. Predictive model

The previous section outlined a method for estimating the function describing a process, done by maximizing the likelihood $P(x, y|v, w)$. Shifting focus to prediction of y given x , the likelihood to maximize is $P(y|x, u)$, where u are the weights of the model f_u . Furthermore, y is no longer allowed as an input to the model, since in this case it would not be predictive. Therefore, the likelihood is expressed as $P(y|x, u) \propto \prod_i \exp\left[-\frac{|y_i - f_u(x_i)|}{2\sigma_y^2}\right]$ and the optimization criterion reduces to the MSE. Conclusively, for prediction, measurement errors of input variables are ignored and the predictive model consists of only one FNN.

III. EXPERIMENT AND DATA

A. Experimental setup

The experimental setup, shown in Fig. 3, consisted of the LUX laser-plasma accelerator with a subsequent electron beamline to generate undulator x-rays. The laser-plasma interaction was driven by the Ti:sapphire laser system Angus, operating at 1 Hz and delivering around 2.6 J pulses with a duration of 34 fs. The laser pulses were focused by a $f/25$ off-axis parabolic mirror (OAP) to a spot size of 25 μm FWHM in the plasma source.

The laser energy was monitored by a pyroelectric sensor ($< 0.5\%$ repeatability) that measured the rejected beam from a polarizer inside a tunable attenuator at the end of the laser chain. The spectral properties of the laser were measured using a fiber-coupled spectrometer (1 nm

resolution) that analyzed the scattering signal from the pyroelectric sensor. The beam profile and wave front properties of the laser were analyzed using the leakage through the off-axis parabolic mirror. The measured wave front data were decomposed into Zernike coefficients for detailed characterization. The focus position of the laser inside the plasma source was inferred from the defocus value of this measurement. The wave front sensor was calibrated relative to the focus position inside the plasma source by shifting a lens to introduce defocus and observing the plasma recombination light at the focus. The repeatability of the wave front sensor (< 4 nm) corresponds to a focus position uncertainty < 100 μm .

The plasma source consisted of a sapphire plate with a capillarylike structure and three gas inlets, providing a region with a mixture of nitrogen and hydrogen in the front

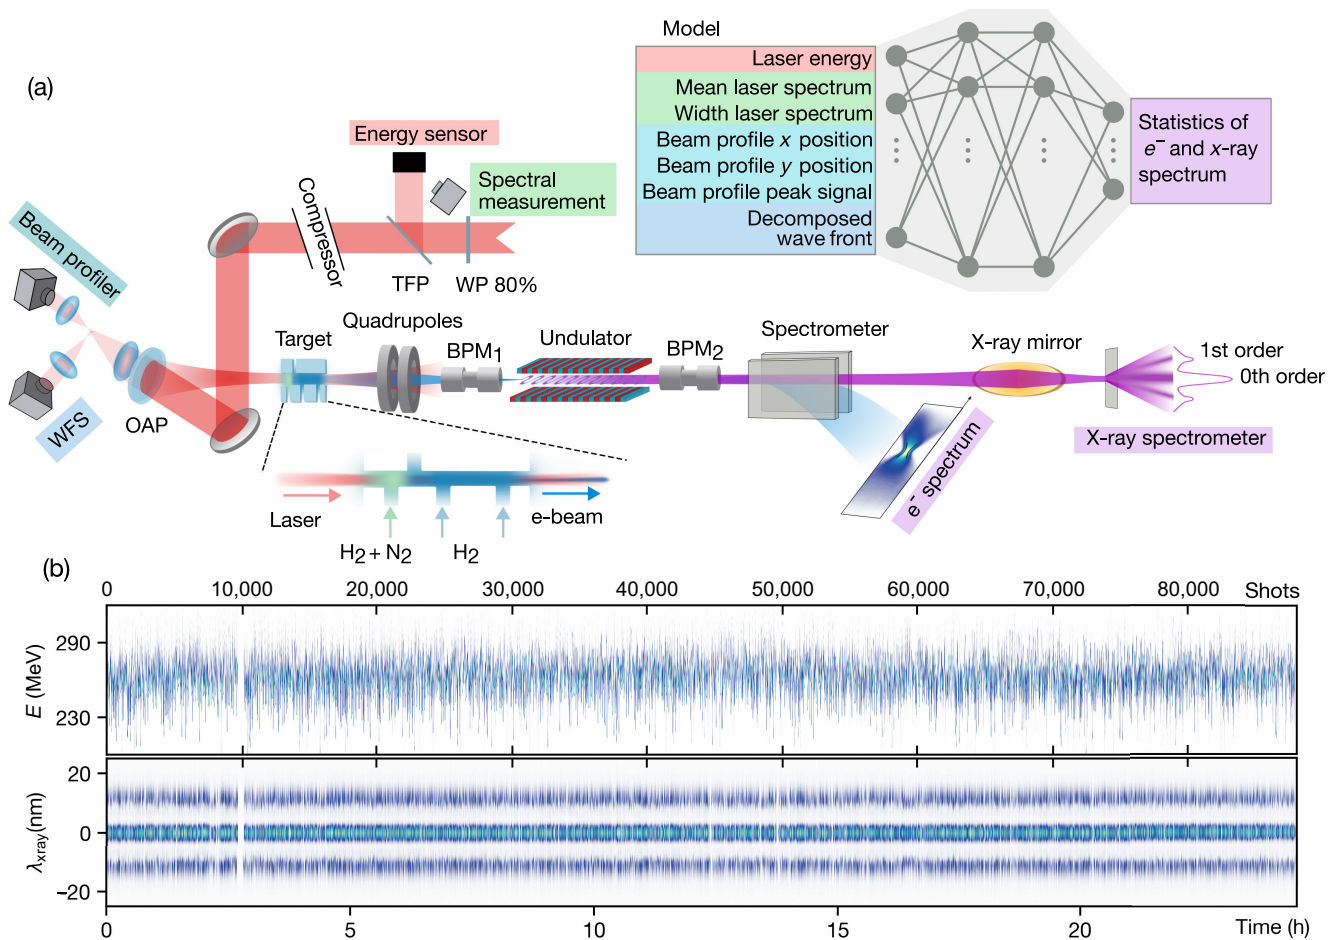


FIG. 3. (a) Shows the experimental setup and (b) the shot-by-shot x-ray and electron spectra over the 24 h long experiment. The laser energy and spectrum were measured before temporarily compressing the laser pulse. Beam profile and wave front were measured from the leakage of the OAP. The OAP focused the laser into the plasma source, where electrons were accelerated using downramp-assisted ionization injection, the electrons were thereafter focused by two quadrupoles into the undulator. The electron spectrum was measured through the deflection of a dipole magnet and the x-rays were focused onto a transmission grating, where the resulting spectrum was captured by a x-ray CCD. Charge and beam position were measured before and after the undulator with beam position monitor (BPM) 1 and 2. The input variables of the explanatory model were measurements of energy, laser spectrum, beam profile, as well as wave front, and output variables were statistics from the electron and x-ray spectrum, as well as the charge measured by BPM₁.

for ionization injection, followed by pure hydrogen region for acceleration. The design enabled electron generation from downramp-assisted ionization injection [11]. This process involves ionizing the outer shell electrons of nitrogen, which creates a density peak within the plasma at the front of the plasma source. After this density peak, the resulting density transition temporarily reduces the wakefield phase velocity, allowing electrons to be locally injected into the plasma wave. These electrons are then accelerated in the density plateau region formed by the pure hydrogen present in the latter part of the plasma source. This setup delivered on average electrons with median energy $E = 260$ MeV (11 MeV RMSD), median absolute deviation $\Delta E = 5.5$ MeV (2.6 MeV RMSD), and charge $Q = 60$ pC (16 pC RMSD).

Electron beams exiting the plasma were focused into the center of the BEAST II undulator by a quadrupole doublet. The undulator, consisting of 100 periods with 5 mm period length, a fixed 2 mm gap, and undulator parameter $K = 0.29$, generated radiation from the accelerated electron beams. Before and after the undulator, the charge and centroid of the electron beam were measured using beam position monitors (BPMs). Following the undulator, the electron beams were directed into a dipole spectrometer, which achieved an energy resolution of 0.1% within a ± 20 MeV range around the central energy.

The spontaneous undulator radiation was focused using a grazing incidence toroidal mirror, which imaged the center of the undulator into a radiation spectrometer. This spectrometer consisted of a free-standing gold grating that dispersed the radiation by wavelength, with the resulting spectrum measured on an x-ray CCD. The x-ray spectrum measured on average had a central wavelength of 12 nm (0.8 nm RMSD) with bandwidth 1.9 nm (0.18 nm RMSD). These diagnostics allowed the simultaneous measurement of the electron and x-ray properties for all individual shots.

The accelerator was operated continuously for more than 24 h, producing more than 80,000 shots. For the analysis, the first 40,000 shots were used, which allowed us to designate a single training set (shots 0–20,000) and a subsequent validation set (shots 20,000–40,000). The split was performed chronologically to ensure that the training period preceded the validation period, thereby avoiding information leakage between the datasets that could compromise the integrity of the validation performance. Using the full dataset would have required multiple training and validation intervals to address long-term laser drifts that necessitate periodic model retraining—a strategy that, while feasible, was set aside for simplicity.

Measurements of distributions and profiles were summarized in statistics, which, together with scalar diagnostics, resulted in 70 defining laser parameters per shot. This extensive dataset allowed for robust analysis of the system's performance and stability, providing a comprehensive foundation for the results presented in the paper.

B. Feature and model parameter selection

Strong correlations between input variables, a phenomenon known as multicollinearity, can obscure the individual contribution of features and reduce model interpretability. To address this, we used the variance inflation factor (VIF) and permutation feature importance [42] to select a subset of informative and independent features.

The VIF is calculated as $VIF = 1/(1 - R_i^2)$, where R_i^2 is the R^2 score of a linear regression model that predicts the i th input variable based on all other input variables. Permutation feature importance quantifies a variable's importance by the decrease in model accuracy when that variable is shuffled. For this, any model with enough expressibility can be used; here we used feedforward neural networks (FNNs). Each FNN consisted of two hidden layers of 32 neurons each. The accuracy of the model was quantified by the root-mean-squared error (RMSE) for the output variables: charge (Q), electron median energy (E), and electron energy median deviation (ΔE).

By demanding VIF below 5 and high importance the set of potential laser diagnostics and statistics used as input variables could be reduced from 70 to 11. For physical interpretability, wave front aberrations with order 9 or higher were excluded from this set. For completeness, all pairs of variables (e.g., beam profile x - and y -positions) were included, even in cases where only one exhibited high importance. This yielded a final input variable set of 14, shown in Fig. 4. In contrast, incorporating all available diagnostics only increased the R^2 score by 3.4%, suggesting that the additional variables offered little new information to the model. The laser focus position was the most critical input variable, causing the RMSE to increase by up to 200% when shuffled, demonstrating its substantial influence on model accuracy. Unexpectedly, the method identified the peak signal of the laser beam profile as an important feature with a low VIF, suggesting that it captures unique information although no direct physical interpretation could be identified.

The output variables of the explanatory model were chosen as Q , E , ΔE as well as x-ray mean wavelength (λ_{xray}) and standard deviation ($\Delta\lambda_{\text{xray}}$).

For the virtual diagnostic, models were trained using different combinations of three input sets: the laser parameters shown in Fig. 4, the electron bunch position and charge measured by the beam position monitors, and the electron energy spectrum, which was reduced to 20 principal components using PCA. The models were trained to predict the x-ray spectra, which were likewise decomposed via PCA and reconstructed during inference using the inverse transform.

In addition to selecting an appropriate set of input variables, there are also a number of model parameters affecting performance, such as number of neurons, number of hidden layers, weighting, and learning rate. A rigorous optimization of these parameters was outside the scope of




FIG. 4. Feature importance for predicting the median energy (E), median energy absolute deviation (ΔE), and charge (Q) of the electron bunch, shown as a percentage increase of RMSE in log scale. Multicollinearity is quantified by the variance inflation factor (VIF). The bar uncertainty marks the 95% quantile of 50 models.

this work; however, the relevant parameters were varied one by one and the loss monitored to find parameters maximizing performance. The explanatory model consists of networks featuring two FNNs with two hidden layers containing 32 neurons each and Tanh activation. For training, an Adam optimizer with a weight decay of $\eta = 10^{-5}$ and a learning rate of 10^{-3} was used. The model was trained for 600 epochs. Details on the optimization of the weighting parameter can be found in Supplemental Material [36]. The predictive model consists of one neural network with two hidden layers containing 64 neurons each with Tanh activation. No weight decay was used since the convergence was already satisfactory. This model was trained for 100 epochs.

The explanatory and predictive model are ensemble models, that is, they consist of 50 networks with identical architecture, each trained on a different dataset created by random sampling of the training data. The final output is given by the median of all networks, see Fig. 2. The models were built using the Python package `Ensemble-PyTorch` [43], modified to accommodate the OD optimization criteria.

C. Particle-in-cell simulations

To compare our approach with conventional LPA modeling, we performed numerical simulations using the quasicylindrical particle-in-cell (PIC) code `FBPIC` [44]. In these simulations, the laser focus position and energy were randomly sampled over ranges of 1.4 J and 2.7 mm, respectively. The plasma profile was derived from computational fluid dynamics simulations. The laser was modeled with a flattened Gaussian beam profile [45] and a Gaussian temporal profile.

The simulations were conducted in a boosted Lorentz frame with $\gamma_{\text{boost}} = 5$ to accelerate computation. A longitudinal resolution of 50 cells per micrometer and a transverse resolution of 3 cells per micrometer were used, with three azimuthal modes to capture the essential physics.

Because uncertainties during the laser characterization and the idealized conditions imposed in the PIC setup

inevitably introduce small systematic offsets, a residual mismatch between experimental inputs and simulated ones is expected. To enable a fair comparison with the explanatory model, the discrete simulation results were interpolated using cubic splines to create a smooth response surface, which was then translated in input space by applying constant offsets in laser focus position and energy. These offsets were optimized by minimizing the difference between function surfaces resulting from simulations and explanatory model. No rescaling was applied, preserving the relative spacing of simulation points and the physical trends predicted by the simulations.

IV. RESULTS

A. Influence of focus position and pulse energy on electron injection and acceleration

To gain insight into the mechanisms governing the electron and subsequent x-ray generation, it is essential to extract the underlying relationships between laser and resulting beam properties. However, the laser parameters were not scanned systematically but varied simultaneously due to the natural shot-to-shot fluctuations in the experiment. Consequently, their individual influence is obscured when simply looking at marginalized representations of the data and any observable correlations are further reduced by noise. To address this, we use the explanatory model that disentangles these coupled effects and helps reveal the underlying correlations otherwise hidden by the multidimensional parameter space.

The effectiveness of this approach is illustrated in Figs. 5(a)–5(c), where the training data is shown as a 2D-histogram and compared with the dependencies of charge, energy, and energy spread on focus position extracted from the explanatory model.

The model shows that shifting the focus position downstream, i.e., to larger values, reduces the beam charge. This is a result of the drop in laser intensity at the upstream region of the plasma source, where the mixed gas is located.




FIG. 5. Left column (a)–(c) shows the explanatory model and PIC simulations as a function of focus position for charge Q , electron energy E , and energy spread ΔE , as well as the distribution of the training data. The middle two columns (d)–(i) show a comparison between modeled surface (left) and PIC simulated (right) in a matched region. The PIC simulated surface (g)–(i) has been extracted from a larger range of simulations, shown in (j)–(l). The simulations, marked by black dots, generate a function surface through a cubic spline. The red rectangle is the region shown in (g)–(i) and is extracted by subtracting the model surface of (d),(f), corresponding to Q and ΔE , while moving over the PIC simulated surface in (j), (l) to find the point of maximal correspondence.

The lower intensity reduces the ionization rate and thereby decreases the injected charge. The opposite trend is found for the electron beam energy, which increases when the focus is shifted downstream. This effect can be attributed to beam loading [46]. In general, the electron beam absorbs energy from the wakefield and thereby suppresses the accelerating field. As fewer electrons are injected (due to lower intensity at the downstream focus), the wakefield remains strong and the beam gains more energy. In contrast, a higher injected charge would extract more energy from the wake, weaken its accelerating field and reducing the final beam energy. Finally, we observe a V-shaped dependence of the energy spread on the focus position, also caused by beam loading. As the electron bunch depletes the wakefield along its length, the accelerating field weakens toward the tail. Near $z_{\text{foc}} = 0$, this depletion balances the natural wakefield gradient so that all longitudinal slices experience similar accelerating fields, yielding a minimal energy spread. Moving away from this optimum alters the field experienced by the bunch tail, increasing the energy spread and giving rise to the V-shaped dependence.

In addition to focus position, the model lets us examine how laser energy affects the beam. As shown in Figs. 5(d)–5(f), electron beam parameters remain largely unaffected by variations in laser energy, in stark contrast to their pronounced dependence on focus position. This

observation is in good agreement with PIC simulations within the same experimental parameter range, see Figs. 5(g)–5(i).

To directly quantify the impact of each of these laser parameters on beam stability, we used the explanatory model to relate shot-to-shot fluctuations of focus position and laser energy to electron energy. For instance, a 1 MeV RMSD fluctuation in electron energy corresponds to a 7 μm RMSD shot-to-shot variation in focus position. To cause the same 1 MeV RMSD variation would require a 0.03 J fluctuation in laser energy. In our experiment, 7 μm represents only 4% of the measured RMSD focus-position shot-to-shot variation, whereas 0.03 J is about 160% of the RMSD laser-energy shot-to-shot variation. Consequently, within the typical operational stability of today's laser systems, the electron beam properties remain far more sensitive to changes in focus position than to variations in laser energy. This finding stands in contrast to the usual emphasis on laser-energy stability in discussions of next-generation systems, suggesting that wave front fluctuations, such as focus-position shot-to-shot variation, may be the more pressing issue for improving beam quality and stability.

Besides validating the explanatory model, the PIC simulations allow exploration of parameter regimes not covered by our measurements. The broader range of laser energy and focus position available through PIC simulations [Figs. 5(j)–5(l)] reveals that laser energy becomes

more influential once varied on a larger scale. As expected, the beam charge increases with higher laser energy and when the focus is closer to the mixed-gas region. Meanwhile, electron energy and energy spread remain primarily dictated by focus position. A plausible explanation is that while focus position predominantly sets the injected charge, laser energy simultaneously affects both the injected charge and the wakefield strength. As a result, a stronger wakefield can sustain a higher charge without overloading, thus leaving the beam energy and energy spread largely unchanged.

Note that the explanatory model and simulations have a maximal overlap at a point where the difference between modeled and simulated laser energy is as low as 5%. This indicates that the model, despite the low significance of laser energy indicated by feature importance, is able to extract a likely dependence between electron properties and laser energy. The overall good agreement between simulation and model, indicates the future possibility of multifidelity modeling, where simulated data is integrated with experimental models to identify optimal points beyond the range of experimental data.

B. Chromatic effects and higher-order wave front aberrations

Additional laser variables of importance, see Fig. 4, included the spectral properties of the laser pulse as well as higher-order wave front aberrations. Variations in the spectral properties have a direct impact on the LPA process. For example, both the laser vector potential and the pulse duration—which are critical for plasma wave excitation—depend on the central wavelength and spectral bandwidth. In our measurements, the RMSD of these parameters was 0.3 and 0.1 nm, respectively. While the variations in bandwidth could result in a percent-level change in laser intensity (via the pulse duration), the direct effects alone cannot fully explain the influence of central wavelength shifts as captured by the model.

Nevertheless, even such small spectral changes can have a substantial indirect influence on the LPA process through *chromatic effects*. A prime example are chromatic aberrations introduced by transmissive optical elements in the beam path, which effectively cause a shift in focus position, $\Delta z_{\text{foc}} = c_{\text{sys}}(\lambda - \bar{\lambda})$, given a change in central wavelength λ relative to the average $\bar{\lambda}$. For our setup, we estimate $c_{\text{sys}} \approx 100 \mu\text{m}/\text{nm}$.

This focus shift should be accounted for by the wave front sensor, as it measures the average focus position of all wavelengths. However, the dependence between laser wavelength, focus position, and electron properties indicates the presence of another correlation between those properties. This can be attributed to additional chromaticity $c_{\text{diag}} \approx 215 \pm 10 \mu\text{m}/\text{nm}$ that is accumulated in the diagnostic beam path toward the wave front sensor (compare




FIG. 6. (a) Shows the model surface for median energy spread (ΔE) as a function of the central laser wavelength deviation $\lambda - \bar{\lambda}$, where $\bar{\lambda}$ is the mean central wavelength over all data samples. (b)–(f) show ΔE as a function of the wave front aberrations: vertical astigmatism, horizontal coma, primary spherical wave front as well as focus position (z_{foc}). The orange dashed line shows the chromaticity of the wave front measurement that is the total chromaticity of the optics in the beam path from the OAP to the wave front diagnostics.

Fig. 3). The induced measurement error $\Delta z_{\text{foc, err}} = c_{\text{diag}}(\lambda - \bar{\lambda})$ effectively leads to equipotential lines on the model surface with slope equal to c_{diag} , as can be seen in Fig. 6(a). For a derivation, we refer to Supplemental Material [36].

We estimate that approximately 30% of the total variation in the measured focus position can be attributed to chromatic effects. In addition, the observation $c_{\text{diag}} \approx 2c_{\text{sys}}$ indicates that measurement errors constitute a significant portion of this variance, rather than reflecting true variations of the focus position in the plasma source. Note that one could, in principle, remove this measurement error of the focus position by subtracting $c_{\text{diag}}(\lambda - \bar{\lambda})$ from z_{foc} , assuming that the spectrum is accurately measured. Nevertheless, a much better approach would be to decrease the chromaticity of the diagnostics in future experiments. In summary, the analysis shows that even though variations in the central laser wavelength have a low direct effect on the acceleration dynamics, they can cause indirect chromatic effects and introduce measurement errors.

The most significant parameters after focus position and central laser wavelength were the higher-order wave front aberrations: vertical astigmatism, horizontal coma, and primary spherical, see Fig. 4. Figure 6 shows that the electron properties have a similar relationship with these parameters as with focus position, namely, along these dimensions there is a minimum of energy spread. This suggests that the higher wave front aberrations mainly affect

the electron spectra through the intensity in the injection region, hence through the process of beam loading. The remaining wave front aberrations, oblique astigmatism and vertical coma, had low variance compared to vertical astigmatism and horizontal coma, which explains their low importance, see Fig. 4.

To gain more knowledge about the influence of wave front aberrations on the LPA, it is necessary to develop datasets with larger variance in the wave front and analyze alongside more realistic 3D PIC simulations than presented in this work. Nevertheless, already at this level of accuracy, it would be possible to use the model surface, partly shown in Fig. 6, for tuning the electron quality by adjusting the wave front.

C. Laser influence on x-ray generation

In the preceding sections, we concentrated on the laser-plasma interaction and its immediate impact on the quality and stability of the electron beam. Now, we shift our focus to the implications these dynamics can have on the generation of secondary radiation.

To build a consistent picture, we begin with a straightforward analysis of the raw experimental data. Figure 7(a) shows the integrated x-ray spectrum sorted by the median electron energy (a moving mean has been applied along this dimension to reduce noise). The spectra exhibit a relatively broad bandwidth, with a clear dependence of the central wavelength λ_{xray} on electron energy E . We can compare this observation with the theoretical relationship between electron energy and x-ray wavelength, which is described by the undulator equation

$$\lambda_{\text{xray}}(\gamma, \theta) = \frac{\lambda_u}{2\gamma^2} \left(1 + \frac{K^2}{2} + \gamma^2\theta^2 \right), \quad (4)$$

where λ_{xray} is the radiation wavelength, λ_u the undulator period, K the undulator parameter, γ the relativistic factor, and θ the observation angle. The LUX beamline has $K = 0.29$ and $\lambda_u = 5$ mm. This relation is shown as a black solid line in Fig. 7(a), assuming on-axis emission ($\theta = 0$).

It is important to note that this theoretical curve intersects the slope of the measured spectra toward shorter wavelengths, rather than coinciding with their peak. This discrepancy is explained by the combined effects of off-axis observation angles, captured within the spectrometer's angular acceptance (~ 6 mrad), and the natural divergence of the electron beam as it propagates through the undulator. Both factors cause a broadening of the radiated spectrum toward longer wavelengths due to the quadratic dependence on θ in the governing Eq. (4). Additionally, transverse variations in the undulator parameter K can lead to further broadening of the spectrum given the finite size of the electron beam. Note that the predictions of the explanatory model [green dashed line in Fig. 7(a)] intrinsically account for these effects, since we use the mean and standard




FIG. 7. (a) Shows a waterfall plot of x-ray spectra, consisting of 40,000 shots, ordered and binned by electron median energy, and normalized by the mean intensity of the peak signal of the 0th order diffraction. Black line marks the analytic estimation for the x-ray wavelength with emission angle $\theta = 0$. (b) Shows the spectra in (a) at three different electron energies marked by red, black, and blue dashed lines in (a). Subplot (c) and (d) show the distributions of median energy and electron energy spread to central wavelength and x-ray bandwidth. The explanatory model in (a), (c), (d) (dashed green) is extracted by scanning focus position, z_{foc} , while fixing all other input to the working point. The focus position is ranging from $z_{\text{foc}} = [-330, 330] \approx [-2\sigma_z, 2\sigma_z]$ μm . Note the difference in strength of correlation between (c) and (d), which is a consequence of the underlying dependence on z_{foc} .

deviation of the measured x-ray spectrum as output parameters.

The spectral shape of the x-ray radiation is further linked to the energy spread of the electron beam. Under optimal beam loading conditions, where the energy spread is minimized ($\Delta E \approx 2$ MeV), broadening effects associated with variations in γ are negligible. However, when the plasma wave is nonideally loaded, deviations in the beam energy occur, accompanied by an increased energy spread. This leads to additional spectral broadening that is particularly visible toward shorter wavelengths.

The behavior is further analyzed in Fig. 7(b), which compares normalized and shifted spectra for three beam loading conditions: underloaded (red), overloaded (blue), and optimally loaded (black). As the energy spread decreases, the lower end of the spectrum steepens, revealing a direct correlation between electron beam quality and x-ray characteristics.

The power of our explanatory model lies in its ability to directly link these observed x-ray spectral effects to specific laser properties. In order to isolate the contribution of individual laser parameters, we can systematically vary them independently and analyze the corresponding model output. Through this approach, we identified that the laser

focus position is particularly decisive in shaping the spectral characteristics of the emitted x-rays.

In Figs. 7(c) and 7(d), we present the measured energy and relative bandwidth of the x-ray and electron spectrum, with the explanatory model as a function of the focus position, z_{foc} . Remarkably, scanning the focus over a range of 660 μm , reproduces the strong correlation between the electron median energy E and the central x-ray wavelength λ_{xray} . As discussed previously, the shift in wavelength originates from changes in the electron energy, which are themselves induced by variations in beam loading.

The same physical principles are evident when comparing the energy spread of the electron beam with the x-ray bandwidth. As the focus is moved toward optimal charge injection, both the electron energy spread and the x-ray bandwidth decrease. However, moving the focus beyond the optimal point reverses this trend, causing both quantities to rise again. Furthermore, there is a difference between x-rays generated by laser pulses with more upstream or downstream focus position, despite having the same relative energy bandwidth. This could be explained by the corresponding change in electron energy, which alters the beam transport. While this effect is hidden in the raw data, the model is able to elucidate the process by extracting the relationship between the focus position and the properties of electrons and x-rays.

Beyond enhancing our understanding of the underlying physics, these insights have practical implications. For instance, the parametrized function could serve as a tuning curve for the x-ray source. Simply adjusting the position of a lens within the final beam expander telescope of the laser system would offer a straightforward method to control the x-ray output. With further refinement and by incorporating additional parameters, such a parametrized model could even enable independent control over multiple beam properties. In the long term, the approach may also form the basis for an active feedback system to maintain stable x-ray output, ensuring consistent performance in future applications.

D. Virtual x-ray diagnostics

The previous sections emphasized the usefulness of the explanatory model to gain insight into the physics. However, in x-ray applications, such as time-resolved pump-probe experiments [47], photoelectron spectroscopy [48], and absorption spectroscopy [49], the primary objective would not be to understand the underlying processes, but simply to obtain accurate knowledge of the radiated spectrum. This becomes particularly important when the source exhibits high variance, as discussed in [22].

A crucial challenge in these applications is that standard diagnostics for characterizing x-ray sources are often destructive, making them unsuitable for real-time use. This is where virtual diagnostics could play a transformative role. By leveraging predictive models, they can provide noninvasive and real-time estimates of the x-ray spectrum,

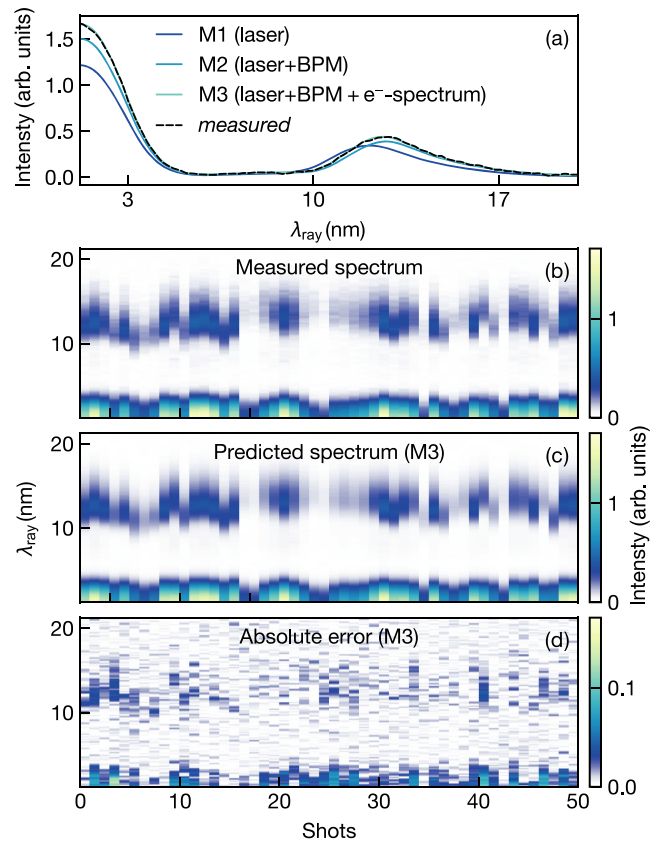


FIG. 8. (a) Shows the prediction for three models with different input variables (M1, M2, M3) and the measured spectrum (black dashed). M1 uses the laser properties shown in Fig. 4, M2 uses the same laser properties and data from two beam position monitors, and M3 uses all aforementioned properties and the measured electron spectra. (b) Shows 50 consecutive measured x-ray spectra and (c) the corresponding prediction of model M3. (d) Shows the residual between measured and predicted spectrum of model M3.

providing precise data without interfering with the experimental setup.

For this predictive task, we compare three models, each using a different set of input variables. Model M1 uses only laser properties, consistent with the input parameters used in the explanatory model. Model M2 adds noninvasive electron diagnostics, specifically data from beam position monitors (BPM), which capture the bunch position and charge. Importantly, this model remains fully noninvasive, as it leaves both the electron beam and the x-ray beam intact. This makes M2 potentially relevant for pump-probe experiments where both beams are used. Finally, model M3 incorporates measurements from the destructive electron spectrometer. This provides additional detailed information about the electron spectrum. M3 is suitable for scenarios where only the x-ray beam is of interest.

Figure 8(a) illustrates the predicted spectrum for a selected shot using each model. While a virtual diagnostic based solely on laser data alone already enables a decent

prediction of the x-ray spectrum (M1), it is outperformed by models that include direct information about the electron beam (M2 and M3), as confirmed by the R^2 scores for predicting the central wavelength, 0.61 (M1), 0.93 (M2), and 0.98 (M3).

Remarkably, merely adding the measurements of the beam charge and position to the input in M2 increases the prediction accuracy by over 50% compared to M1. This gain reflects the fact that these measurements provide direct information on key beam properties: the charge encodes aspects of the energy spectrum via beam loading, while the position reflects the trajectory through the beamline, which not only determines how the beam interacts with the undulator field but also carries information about the beam energy. The importance of this additional information is even more pronounced in predicting the x-ray flux (integrated first order signal), with R^2 scores of 0.33 (M1), 0.91 (M2), and 0.97 (M3). The limited accuracy of M1 indicates that the laser parameters alone may not sufficiently capture the initial beam pointing and resulting trajectory, which in turn provide indirect information on possible losses during transport.

Despite relying only on noninvasive diagnostics, M2 achieves outstanding accuracy and is highly suitable for applications where destructive electron beam diagnostics are impractical. Still, additionally incorporating the full electron spectrum in model M3 delivers the best performance, as further demonstrated by the close agreement between the measured and predicted spectra in Figs. 8(b) and 8(c). A comprehensive list of model performance with all combinations of input parameters can be found in Supplemental Material [36].

V. CONCLUSION

In summary, we have shown that data-driven models, when carefully designed, can capture the complex interplay between laser parameters and electron dynamics in a laser-plasma accelerator. Central to our approach is a model architecture that leverages an orthogonal-distance-based training objective to effectively minimize the impact of measurement errors in input parameters. The model effectively uncovers the dependencies between laser parameters and electrons, eventually enabling tracing variations in the resulting x-ray spectra back to specific changes in the drive laser. Moreover, our predictive model precisely forecasts the x-ray spectrum, offering a noninvasive diagnostic that can provide prospective users with shot-to-shot parameters of the delivered radiation. The same model framework could also serve as the basis for active feedback loops that directly stabilize the secondary radiation during operation.

Recent advances have highlighted the potential of laser-plasma accelerators as compact drivers of secondary radiation sources. However, their stability, reliability, and tunability still fall short of what is routinely achieved with conventional accelerator technologies. Our approach helps

to address this gap by guiding the development of more stable LPA systems and enabling model-based control and feedback strategies.

Looking ahead, the increase in high-performance computing resources and experimental data—fueled by the upcoming era of high-repetition-rate lasers—will only increase the relevance of data-driven modeling in the field. Leveraging these developments will be critical to transforming laser-plasma accelerators from proof-of-concept systems into practical, application-ready machines.

ACKNOWLEDGMENTS

This work was supported by the Swedish Research Council Grant No. 2023-06998. Computational resources were provided by the Maxwell HPC cluster operated at Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

-
- [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. Picksley, J. Stackhouse, C. Benedetti, K. Nakamura, H. E. Tsai, R. Li, B. Miao, J. E. Shrock, E. Rockafellow, H. M. Milchberg, C. B. Schroeder, J. van Tilborg, E. Esarey, C. G. R. Geddes, and A. J. Gonsalves, Matched guiding and controlled injection in dark-current-free, 10-GeV-class, channel-guided laser-plasma accelerators, *Phys. Rev. Lett.* **133**, 255001 (2024).
 - [4] O. Lundh, J. Lim, C. Rechatin, L. Ammoura, A. Ben-Ismaïl, X. Davoine, G. Gallot, J.-P. Goddet, E. Lefebvre, V. Malka, and J. Faure, Few femtosecond, few kiloampere electron bunch produced by a laser-plasma accelerator, *Nat. Phys.* **7**, 219 (2011).
 - [5] Z. Guo, S. Liu, B. Zhou, J. Liu, H. Wang, Y. Pi, X. Wang, Y. Mo, B. Guo, J. Hua, Y. Wan, and W. Lu, Preclinical tumor control with a laser-accelerated high-energy electron radiotherapy prototype, *Nat. Commun.* **16**, 1895 (2025).
 - [6] J. Wenz, S. Schleede, K. Khrennikov, M. Bech, P. Thibault, M. Heigoldt, F. Pfeiffer, and S. Karsch, Quantitative x-ray phase-contrast microtomography from a compact laser-driven betatron source, *Nat. Commun.* **6**, 7568 (2015).
 - [7] W. Wang, K. Feng, L. Ke, C. Yu, Y. Xu, R. Qi, Y. Chen, Z. Qin, Z. Zhang, M. Fang *et al.*, Free-electron lasing at 27 nanometres based on a laser wakefield accelerator, *Nature (London)* **595**, 516 (2021).

- [8] M. Labat *et al.*, Seeded free-electron laser driven by a compact laser-plasma accelerator, *Nat. Photonics* **17**, 150 (2023).
- [9] S. K. Barber, F. Kohrell, C. E. Doss, K. Jensen, C. Berger, F. Isono, Z. Eisentraut, S. Schröder, A. J. Gonsalves, K. Nakamura, G. R. Plateau, R. A. van Mourik, M. Gracia-Linares, L. Labun, B. M. Hegelich, S. V. Milton, C. G. R. Geddes, J. Osterhoff, C. B. Schroeder, E. H. Esarey, and J. van Tilborg, Greater than 1000-fold gain in a free-electron laser driven by a laser-plasma accelerator with high reliability, *Phys. Rev. Lett.* **135**, 055001 (2025).
- [10] H. T. Kim, V. B. Pathak, K. Hong Pae, A. Lifschitz, F. Sylla, J. H. Shin, C. Hojbota, S. K. Lee, J. H. Sung, H. W. Lee, E. Guillaume, C. Thaury, K. Nakajima, J. Vieira, L. O. Silva, V. Malka, and C. H. Nam, Stable multi-GeV electron accelerator driven by waveform-controlled PW laser pulses, *Sci. Rep.* **7**, 10203 (2017).
- [11] 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).
- [12] A. R. Maier, N. M. Delbos, T. Eichner, L. Hübner, S. Jalas, L. Jeppe, 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).
- [13] L. Rovige, J. Huijts, I. Andriyash, A. Vernier, V. Tomkus, V. Girdauskas, G. Raciukaitis, J. Dudutis, V. Stankevicius, P. Gecys, M. Ouille, Z. Cheng, R. Lopez-Martens, and J. Faure, Demonstration of stable long-term operation of a kilohertz laser-plasma accelerator, *Phys. Rev. Accel. Beams* **23**, 093401 (2020).
- [14] A. Döpp, C. Eberle, S. Howard, F. Irshad, J. Lin, and M. Streeter, Data-driven science and machine learning methods in laser-plasma physics, *High Power Laser Sci. Eng.* **11**, e55 (2023).
- [15] M. Streeter *et al.*, Laser wakefield accelerator modelling with variational neural networks, *High Power Laser Sci. Eng.* **11**, e9 (2023).
- [16] S. J. D. Dann *et al.*, Laser wakefield acceleration with active feedback at 5 Hz, *Phys. Rev. Accel. Beams* **22**, 041303 (2019).
- [17] A. Hanuka, C. Emma, T. Maxwell, A. S. Fisher, B. Jacobson, M. J. Hogan, and Z. Huang, Accurate and confident prediction of electron beam longitudinal properties using spectral virtual diagnostics, *Sci. Rep.* **11**, 2945 (2021).
- [18] 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).
- [19] R. J. Shalloo *et al.*, Automation and control of laser wakefield accelerators using Bayesian optimization, *Nat. Commun.* **11**, 6355 (2020).
- [20] J. Lin, Q. Qian, J. Murphy, A. Hsu, A. Hero, Y. Ma, A. G. R. Thomas, and K. Krushelnick, Beyond optimization—supervised learning applications in relativistic laser-plasma experiments, *Phys. Plasmas* **28**, 083102 (2021).
- [21] O. Convery, L. Smith, Y. Gal, and A. Hanuka, Uncertainty quantification for virtual diagnostic of particle accelerators, *Phys. Rev. Accel. Beams* **24**, 074602 (2021).
- [22] A. Sanchez-Gonzalez *et al.*, Accurate prediction of x-ray pulse properties from a free-electron laser using machine learning, *Nat. Commun.* **8**, 15461 (2017).
- [23] C. Emma, A. Edelen, M. Hogan, B. O’Shea, G. White, and V. Yakimenko, Machine learning-based longitudinal phase space prediction of particle accelerators, *Phys. Rev. Accel. Beams* **21**, 112802 (2018).
- [24] G. James, D. Witten, T. Hastie, R. Tibshirani, and J. Taylor, Statistical learning, in *An Introduction to Statistical Learning: With Applications in Python* (Springer International Publishing, Cham, 2023), p. 17, 10.1007/978-3-031-38747-0.
- [25] G. Shmueli, To explain or to predict?, *Stat. Sci.* **25**, 289 (2010).
- [26] P. Mehta, M. Bukov, C. Wang, A. G. R. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, A high-bias, low-variance introduction to machine learning for physicists, *Phys. Rep.* **810**, 1 (2019).
- [27] R. J. Carroll, D. Ruppert, L. A. Stefanski, and C. M. Crainiceanu, *Measurement Error in Nonlinear Models*, 2nd ed. (Chapman and Hall/CRC, London, 2006), p. 1, 10.1201/9781420010138.
- [28] J. A. Hutcheon, A. Chiolero, and J. A. Hanley, Random measurement error and regression dilution bias, *Br. Med. J.* **340**, c2289 (2010).
- [29] R. J. Adcock, A problem in least squares, *Analyst* **5**, 53 (1878).
- [30] C. H. Kummell, Reduction of observation equations which contain more than one observed quantity, *Analyst* **6**, 97 (1879).
- [31] L. Leng, T. Zhang, L. Kleinman, and W. Zhu, Ordinary least square regression, orthogonal regression, geometric mean regression and their applications in aerosol science, *J. Phys. Conf. Ser.* **78**, 012084 (2007).
- [32] J. A. Hausman, W. K. Newey, H. Ichimura, and J. L. Powell, Identification and estimation of polynomial errors-in-variables models, *J. Econom.* **50**, 273 (1991).
- [33] S. Schennach, Estimation of nonlinear models with measurement error, *Econometrica* **72**, 33 (2004).
- [34] Z. Hu, Z. T. Ke, and J. S. Liu, Measurement error models: From nonparametric methods to deep neural networks, *Stat. Sci.* **37**, 473 (2022).
- [35] K. Hornik, Approximation capabilities of multilayer feed-forward networks, *Neural Netw.* **4**, 251 (1991).
- [36] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/dnnw-lfn5> for details on choice of model parameters σ_x, σ_y and weight regularization in the orthogonal distance model as well as derivation of effects of chromaticity induced measurement errors and additional information on the predictive model.
- [37] K. P. Murphy, Machine learning: A probabilistic perspective, in *Adaptive Computation and Machine Learning Series* (MIT Press, Cambridge, MA, 2012), pp. 225–227, 252.
- [38] S. C. Hora, Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management, *Reliab. Eng. Syst. Saf.* **54**, 217 (1996).
- [39] M. H. Shaker and E. Hüllermeier, Aleatoric and epistemic uncertainty with random forests, in *Advances in Intelligent Data Analysis XVIII* (Springer International Publishing, 2020), p. 444, 10.1007/978-3-030-44584-3_35.

- [40] B. Efron, Bootstrap methods: Another look at the jackknife, *Ann. Stat.* **7**, 1 (1979).
- [41] J. Franke and M. H. Neumann, Bootstrapping neural networks, *Neural Comput.* **12**, 1929 (2000).
- [42] C. Molnar, 8.5 permutation feature importance, in *Interpretable Machine Learning* (Leanpub, 2020), <https://christophm.github.io/interpretable-ml-book/feature-importance.html>.
- [43] Ensemble-PyTorch, PyTorch (software), version 0.9.1, <https://github.com/TorchEnsemble-Community/Ensemble-Pytorch>.
- [44] 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).
- [45] M. Santarsiero, D. Aiello, R. Borghi, and S. Vicalvi, Focusing of axially symmetric flattened Gaussian beams, *J. Mod. Opt.* **44**, 633 (1997).
- [46] 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).
- [47] A. Stolow, A. E. Bragg, and D. M. Neumark, Femtosecond time-resolved photoelectron spectroscopy, *Chem. Rev.* **104**, 1719 (2004).
- [48] J. Chastain and R. C. King, Jr., *Handbook of X-Ray Photoelectron Spectroscopy*, edited by J. Chastain and R. C. King, Jr (Perkin-Elmer Corporation, Eden Prairie, MN, 1992), Vol. 40, p. 10.
- [49] J. Yano and V. K. Yachandra, X-ray absorption spectroscopy, *Photosynth. Res.* **102**, 241 (2009).