

<https://doi.org/10.1038/s42005-026-02542-6>

# Batch Bayesian optimization of attosecond betatron pulses from laser wakefield acceleration

Check for updates

Dominika Maslarova <sup>1</sup>✉, Albert Hansson <sup>1</sup>, Mufei Luo <sup>1,2</sup>, Vojtěch Horný <sup>1,3</sup>, Julien Ferri <sup>1</sup>, Istvan Pusztai <sup>1</sup> & Tünde Fülöp <sup>1</sup>

Laser wakefield acceleration can generate a femtosecond-scale broadband X-ray betatron radiation pulse from electrons accelerated by an intense laser pulse in a plasma. The micrometer-scale of the source makes wakefield betatron radiation well-suited for advanced imaging techniques, including diffraction and phase-contrast imaging. Recent progress in laser technology can expand these capabilities into the attosecond regime, where the practical applications would significantly benefit from the increased energy contained within the pulse. Here we use numerical simulations combined with batch Bayesian optimization to enhance the radiation produced by an attosecond betatron source. The method enables an efficient exploration of a multi-parameter space and identifies a regime in which a plasma density spike triggers the generation of a high-charge electron beam. This results in an improvement of more than one order of magnitude in the on-axis time-averaged power within the central time containing half of the radiated energy, compared to the reference case without the density spike.

The generation of electromagnetic pulses on attosecond timescales represents a milestone in ultrafast science. These pulses offer unprecedented time resolution, enabling real-time observation of electron dynamics in atoms and molecules. A major advancement in attosecond pulse generation has been achieved through high-order harmonic generation<sup>1–4</sup>, where intense laser fields produce extreme ultraviolet and soft X-ray radiation. Attosecond pulses can also be generated with femtosecond-scale intense laser pulses through electron acceleration and associated betatron radiation, as demonstrated by numerical simulations<sup>5,6</sup>. Betatron radiation is produced during laser wakefield acceleration (LWFA)<sup>7</sup>, a laser-plasma interaction technique for accelerating electrons to relativistic speeds. While LWFA electron beams typically have durations on femtosecond timescales—corresponding to the typical dimensions of the wakefield structure—it is also possible to push the beam durations into the sub-fs regime<sup>5,6,8–15</sup>, with betatron radiation inheriting the attosecond scale<sup>5,6</sup>.

LWFA betatron radiation, characterized by broadband X-ray spectra<sup>16–18</sup>, offers good spatial coherence, as demonstrated in femtosecond-scale X-ray phase contrast imaging of biological<sup>19–21</sup> and complex microstructure<sup>22</sup> samples. Extending photon energies of attosecond betatron pulses into the keV range enables X-ray absorption spectroscopy of high-Z elements<sup>23</sup>, and simultaneous multi-element probing in

single-shot measurements<sup>24</sup>. Attosecond betatron radiation thus has the potential to significantly complement techniques like high-harmonic generation to enhance temporal resolution and spatial coherence, improving the ability to focus X-ray beams for high-resolution imaging and precise measurements in ultrafast science.

In LWFA, an ultrashort, ultraintense laser pulse travels through an underdense plasma medium, driving a plasma wave (wakefield) that creates an accelerating gradient that is of the order of a thousand times higher than that of conventional radiofrequency accelerators. The wakefield wavelength is on the order of the plasma wavelength,  $\lambda_p = 2\pi c \sqrt{\epsilon_0 m_e / (e^2 n)}$ , where  $c$  is the speed of light,  $\epsilon_0$  is the vacuum permittivity,  $m_e$  is the electron mass,  $e$  is the elementary charge and  $n$  is the electron density. For laser intensities corresponding to normalized vector potential  $a_0 = 0.855 \lambda_L [\mu\text{m}] \sqrt{I_L [10^{18} \text{W cm}^{-2}]} \gtrsim 2$ , where  $I_L$  is the laser intensity, and  $\lambda_L$  is the laser wavelength, the radiation pressure of the laser pulse expels plasma electrons radially outward, creating ion cavities surrounded by electron sheaths. The first cavity following the laser pulse, the “bubble”, offers the strongest acceleration. This highly nonlinear regime of acceleration is therefore also referred to as the bubble regime. The accelerating field is located at the rear half of the bubble with respect to the pulse propagation direction. During the acceleration in the bubble, the electron beam

<sup>1</sup>Department of Physics, Chalmers University of Technology, Göteborg, Sweden. <sup>2</sup>Department of Physics, University of Oxford, Oxford, UK. <sup>3</sup>Extreme Light Infrastructure - Nuclear Physics, “Horia Hulubei” National Institute for R&D in Physics and Nuclear Engineering, Măgurele, Romania.

✉ e-mail: [dommas@chalmers.se](mailto:dommas@chalmers.se)

undergoes transverse betatron oscillations due to the presence of transverse focusing forces, making the bubble also serve as a wiggler. This mimics the principle of synchrotron devices<sup>25,26</sup>, where the radiated power generally increases with particle energy and the curvature of the trajectory. In the case of the betatron radiation in the bubble regime<sup>25,26</sup>, both the number of photons and the critical energy—below which half of the radiation power is emitted—increase with the charge of the electron beam, electron energy and oscillation amplitude<sup>27</sup>.

Practical applications generally benefit from increased energy contained within the pulse, as this leads to a stronger detectable signal in imaging. Several methods to improve the properties of the betatron radiation have been demonstrated with femtosecond electron beams<sup>17,18,28–44</sup>. For instance, tailoring the plasma density through transverse density gradients<sup>40,41</sup>, longitudinal density shaping<sup>42,43</sup>, or additional high-density stages<sup>44</sup> can enhance betatron emission. A density increase along the propagation direction shortens the plasma wavelength and contracts the bubble, shifting electrons toward the rear of the cavity where the accelerating fields intensify, enabling a phase-reset that renews their energy gain<sup>45</sup> and affects the properties of the betatron radiation<sup>42</sup>.

It is generally challenging to optimize several laser-plasma parameters that can affect the properties of betatron radiation at once due to the high computational cost of kinetic simulations. Bayesian optimization offers an efficient solution by using a probabilistic surrogate model to capture the relationship between adjustable input parameters of the laser-plasma system, and output parameters quantifying the properties of the electron beam and the betatron radiation. In LWFA, it has been applied to improve electron beams in experiments<sup>46–49</sup> and numerical simulations<sup>47,48,50–53</sup>, as well as to enhance betatron-source performance<sup>46,54</sup>. These studies have shown that optimizing parameters such as laser focal position, plasma density, laser spectral phase and plasma length can substantially increase betatron photon yield and improve X ray imaging quality<sup>46,54</sup>.

In this study, we employ batch Bayesian optimization (BBO) to improve the betatron source on an attosecond scale, using three-dimensional (3D) particle-in-cell (PIC) simulations. BBO is an extension of Bayesian optimization that simultaneously selects multiple input points in one iteration, as opposed to evaluating one point at a time, which is practical for a parallelization of the data acquisition when the simulations are time-consuming. Using a density spike in the plasma structure and optimizing its position and gradients, we observe more than 25 times enhancement of the peak of the on-axis radiated energy and more than 6 times increase of the energy contained within the central 50% of the betatron pulse, with only a few tens of PIC simulations requiring a few thousand core hours each. The simulations demonstrate concentrated betatron radiation generation on a millijoule-class laser system using a plasma target, highlighting the potential of low-cost, compact setups for producing attosecond X-ray pulses.

## Results

### Enhancement of the betatron radiation

In our simulations, the attosecond electron beam is produced and injected into the wakefield by a down-ramp injection method<sup>55</sup>, where a density gradient is applied at the plasma entrance. The gradient changes the wake phase velocity, leading to wave breaking, where electrons outrun the rear part of the bubble and consequently are injected into it, where they can be further accelerated. The betatron enhancement was expected to occur due to a locally increased plasma density after a distance  $d_w$ , which shifts electrons to the rear part of the bubble, allowing them to gain more energy during the “phase-reset” mechanism<sup>45</sup>.

We optimize the distance between the injection of the electrons and the start of this density spike  $d_w$ , the length of the spike  $d_s$ , and its peak density  $n_p$ , as shown in Fig. 1. The optimization results indicate that the central value of the betatron energy increases significantly, with a simultaneous reduction in pulse duration when the spike peak density is four times higher than the reference uniform density, the density spike is located shortly (few  $\mu\text{m}$ ) after the injection gradient at the entrance and the spike length is about 120  $\mu\text{m}$ .

The enhancement is actually driven by a second intense electron injection, which produces the significantly more pronounced radiation peak compared to the reference, no-spike configuration.

### Attosecond pulse generation

We consider the interaction of a laser pulse, propagating in the  $x$ -direction, with an underdense plasma. The laser has a wavelength  $\lambda_L = 800$  nm and peak intensity  $2.53 \times 10^{19}$  W cm<sup>-2</sup>, corresponding to a normalized vector potential  $a_0 = 3.44$  and pulse energy  $E_L = 37$  mJ. It has a Gaussian spatial and temporal profile, with linear polarization along the  $y$ -axis. The pulse is focused at the plasma entrance to a spot size of 3.8  $\mu\text{m}$  and has a duration of 8.3 fs (full width at half maximum (FWHM) of the intensity). The set-up is illustrated in Fig. 1. At the entrance of the plasma, a region with longitudinal density variation of total length  $L_g = L_{g_1} + L_{g_2} = 40$   $\mu\text{m}$  is introduced to trigger the initial electron injection. The plasma density rises linearly from zero to  $1.2n_0$ , with  $n_0 = 2 \times 10^{19}$  cm<sup>-3</sup>, over the first  $L_{g_1} = 30$   $\mu\text{m}$ , and then decreases back to  $n_0$  over the following  $L_{g_2} = 10$   $\mu\text{m}$ . This configuration was also used by Ferri et al.<sup>6</sup>, to numerically demonstrate the generation of an attosecond betatron pulse from a down-ramp injection in LWFA.

To enhance the radiation output, we introduce an additional density spike starting at the position  $x = L_g + d_w$ . For simplicity, the linear up and down ramps of the spike are chosen to be symmetric, with a density reaching its maximum value  $n = n_p$  in the middle of the  $d_s$  long region. In the following, the set of values  $\{d_w, d_s, n_p\}$  varied during the optimization is referred to as a point in the parameter space. We set the total plasma length in the  $x$  direction to the dephasing length  $L_d = 2\lambda_{p0}^3 \sqrt{a_0} / (3\pi\lambda_L^2) = 256$   $\mu\text{m}$ , which is the distance at which the electrons outrun the accelerating wakefield region<sup>56</sup>. Here,  $\lambda_{p0}$  is the plasma wavelength at the density  $n_0$ . Despite non-uniform density in the simulations,  $L_d$  reliably estimates the acceleration length  $L_{acc}$ .

To simulate the production of attosecond pulses, we employ a combination of two codes. The interaction of the laser pulse and the plasma is modeled using 3D PIC simulations with Smilei<sup>57</sup>. Based on the electron trajectories obtained from Smilei, the betatron radiation is computed using the Far-field Intensity through Kinetic Analysis (FIKA) code<sup>58</sup>, a Python-based computational tool we developed to analyze the far-field radiation profile from accelerated particles.


### Optimization of betatron radiation

The goal of the optimization is to maximize the radiated energy of the attosecond betatron pulse, which is primarily directed along the pulse propagation direction, i.e., the  $x$  axis. For this reason, we aim to increase the value of  $\frac{d^2 W}{d\Omega dt} |_{\text{on-axis}}$ , which is the on-axis radiated energy  $W$  per time  $t$  per solid angle  $\Omega$ . At the same time, we must ensure that the duration of the betatron pulse remains in the attosecond range, similarly to our reference case that has no density spike. To capture this trade-off, the cost function, aimed to be minimized, is therefore designed as follows:


$$C(W_{50}, \tau_{50}) = -\frac{W_{50}}{\tau_{50}}, \quad (1)$$

where  $W_{50}$  is the energy per solid angle contained within the central 50% of  $\frac{d^2 W}{d\Omega dt} |_{\text{on-axis}}$ , i.e., between the 25% and 75% of the cumulative curve, and  $\tau_{50}$  is the time interval over which  $W_{50}$  is emitted. The “betatron pulse” is defined as the portion of the signal that ends once the  $\frac{d^2 W}{d\Omega dt} |_{\text{on-axis}}$ , after reaching its peak, stays below 1% of that peak for at least 5 as. With this choice of the cost function, longer pulse durations are penalized. Constructing the cost function as a ratio is advantageous, as then a (potentially arbitrary) normalization of the individual figures of merit is not required. Another advantage is that the cost function, due to its integral form, remains stable even when the raw radiation signal contains small oscillations, and therefore does not require any additional smoothing of the radiation profile.

Note that in the reference case additional, weaker betatron pulses appear at later observation times. In our simulation setup we observe two such secondary peaks. To verify their physical origin and relative



**Fig. 1 | An illustration of the proposed setup for betatron radiation enhancement.** The plasma density profile (gray thick line) varies along the longitudinal coordinate  $x$ , while remaining constant in the  $y$  and  $z$  coordinates. The laser pulse first enters a gradient of length  $L_g = 40 \mu\text{m}$ , providing the injection of the electron ( $e^-$ ) beam into the bubble. The  $e^-$  beam produces a betatron pulse that is enhanced by a density spike introduced further along  $x$ . In the simulations, the length of the density spike  $d_s$ , the length of the uniform region with density  $n_0$  before the spike  $d_u$  and the maximum density value of the density spike  $n_p$  were adjusted in the optimization process. The total plasma length is equal to  $L_{acc} = 256 \mu\text{m}$ .



**Fig. 2 | Convergence and wall-clock time of BBO for different batch sizes.** **a** Evolution of the best-performing normalized cost value  $C/C^{\text{ref}}$  achieved up to corresponding iteration for different batch sizes:  $N = 1$  (solid blue),  $N = 4$  (dashed yellow),  $N = 8$  (dotted green). The dash-dotted magenta line marks the reference baseline, i.e., the case with  $C/C^{\text{ref}} = 1$ . **b** Accumulated computational wall-clock time versus iteration for different batch sizes.

contribution, we performed an extended reference simulation covering eight wakefield periods, in which betatron radiation is generated during each plasma period. In the later periods, however, the accelerating fields are significantly weaker, resulting in electron beams with lower charge and energy. Consequently, the combined radiation from the extra seven periods amounts to only around 14% of the betatron on-axis energy. The cost function in Eq. (1) is defined to focus on improving the main first, strongest betatron pulse, which is the pulse of primary interest in this study. As a consequence, the secondary pulses are neither rewarded nor penalized. Optimizing these later pulses would require a different metric depending on a specific application aim, and is beyond the scope of the present work.

For each iteration of the BBO, we performed a simulation batch consisting of a set of  $N$  Smilei electron-trajectory simulations followed by FIKA radiation calculations. One batch corresponds to a complete Smilei-FIKA pipeline (see Methods section and Supplementary Note 1 for more details on the optimization process). Unlike the sequential approach, where the cost function model is updated after each individual evaluation, BBO

updates the model only after a full batch has been evaluated. This allowed us to run multiple simulations in parallel and simultaneously obtain the values of  $W_{50}$  and  $\tau_{50}$  from all of them. After each iteration, the BBO model for each  $N$  was updated manually with the new  $W_{50}$  and  $\tau_{50}$  values, and a new batch of simulations was run. This update step required only a few minutes and the time was comparable for each  $N$ .

BBO provided the flexibility to manually inspect the calculations after the evaluation, while avoiding the long wait times typically associated with updating the model after every evaluation. This approach also enabled close monitoring of the optimization process. This is particularly relevant for strongly nonlinear processes, where it is difficult a priori to know that limiting the search space in a specific way will not lead to solutions of significantly different character – posing different resolution requirements, being physically uninteresting or poorly quantified by the cost function.

To limit the search space to relevant regions, we imposed the following parameter restrictions:

$$d_u \in [0, L_{acc} - L_g - d_s^{\min}], \quad (2)$$

$$d_s \in [d_s^{\min}, L_{acc} - L_g], \quad (3)$$

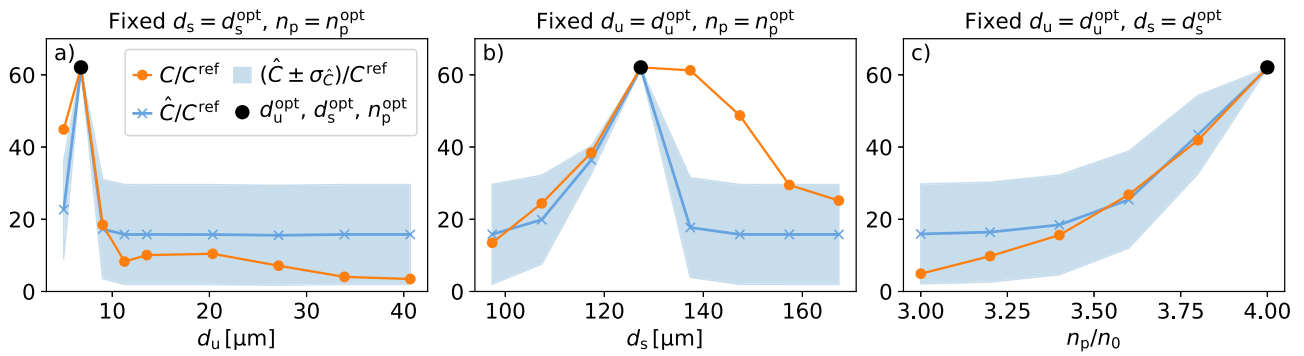
$$n_p \in (n_0, 4n_0), \quad (4)$$

$$d_u + d_s \leq L_{acc} - L_g. \quad (5)$$


Here, we set the minimum spike length to  $d_s^{\min} = 5 \mu\text{m}$ . This small value was arbitrarily chosen to avoid exploring the density profile with  $d_s = 0$ , which corresponds to the reference case. For  $n_p$ , we selected densities greater than  $n_0$  with an upper threshold of  $4n_0$ . This value corresponds to the density where the wakefield wavelength shrinks to  $\lambda_{p0}/2$ , matching the acceleration length in the bubble (the bubble radius) in the plasma density of  $n_0$ . The condition in Eq. (5) ensures that the combined length of the constant density and spike regions does not exceed the total acceleration length. This guarantees that the entire density spike fits within the acceleration length. Because the search space allows for many reasonable choices of parameter bounds, we fixed the bounds a priori and did not optimize them in advance. The results, therefore, reflect improvements within this specified domain. Alternative or expanded limits could potentially lead to other enhancement schemes—for example, higher plasma density for the accelerator-radiator arrangement<sup>44</sup> or longer plasma favoring additional interaction of the electrons with the laser pulse<sup>17,18,28,29</sup>. While such alternative configurations might yield additional improvements, analyzing these dependencies lies beyond the present scope of this work.

To initiate the optimization process, we first simulated the reference case without the spike, i.e., with  $d_s = 0$ , and obtained the reference values of  $W_{50}$  and  $\tau_{50}$ :  $W_{50}^{\text{ref}} = 49.3 \text{ nJ sr}^{-1}$  and  $\tau_{50}^{\text{ref}} = 437 \text{ as}$ , respectively. From these values, we calculated the reference cost from Eq. (1) as  $C^{\text{ref}} = -112.8 \text{ MW sr}^{-1}$ . Subsequently, in the zeroth iteration, we selected  $N_0 = 8$  initial points using a Sobol pseudo-random generator. Sobol sampling provides more uniform coverage of the three-dimensional search space than random sampling, avoiding clustering and improving exploration in the early stages. We then performed eight optimization iterations using batch sizes of  $N = 1, 4$ , and  $8$  (for a total of  $N_0 + N \times 8$  simulations) in order to compare their convergence and performance. The case  $N = 1$  corresponds to the conventional sequential Bayesian optimization scheme.

The cost function  $C$  corresponding to the best-so-far performing case for each iteration (i.e., the lowest value of Eq. (1) found up to the corresponding iteration) are shown in Fig. 2a, for the batch sizes of  $N = 1, 4$ , and  $8$ . Among the tested batch sizes,  $N = 4$  shows the fastest improvement within the eight performed iterations. Its best value,  $C_{N=4}^{\text{opt}} = 60.0 C^{\text{ref}}$  was reached already in the 5th iteration. After that, no further improvement was observed. For  $N = 8$ , the best-performing case appeared later, in the 7th iteration, and was only slightly higher than for  $N = 4$ :  $C_{N=8}^{\text{opt}} = 62.1 C^{\text{ref}}$ .



**Fig. 3 | Values of the normalized cost function around the best-performing point found by BBO.** The values are shown for a) varying  $d_u$ , fixed  $d_s = d_s^{\text{opt}}$ ,  $n_p = n_p^{\text{opt}}$ , b) varying  $d_s$ , fixed  $d_u = d_u^{\text{opt}}$ ,  $n_p = n_p^{\text{opt}}$ , and c) varying  $n_p/n_0$ , fixed  $d_u = d_u^{\text{opt}}$ ,  $d_s = d_s^{\text{opt}}$ . The normalized cost function  $C/C^{\text{ref}}$  is shown with orange lines and circular markers at the evaluated points. The normalized model prediction  $\hat{C}/C^{\text{ref}}$  for  $N = 8$  from the final (8th) iteration, is shown with blue lines with crosses at evaluated points. Colored shaded regions indicate the area within predicted standard deviation  $\sigma_{\hat{C}}$ . The large black circle marks the best-performing case.



**Fig. 4 | Betatron radiation characteristics for the reference and best-performing cases.** **a** On-axis radiated energy  $W$  per time  $t$  per solid angle  $\Omega$  (temporal profile) as a function of the observer's time  $t$  for the best-performing optimized (red dotted) and reference (blue solid) cases. The first two insets show zoomed-in profiles of the betatron peaks in **b**) the best-performing and **c**) reference cases, corresponding to the shaded areas in **a**. **d** On-axis radiated energy  $W$  per photon frequency  $\omega$  per solid angle  $\Omega$  of the betatron pulses as a function of the photon energy  $E_{\text{ph}}$ . Critical energies for the reference case  $E_c^{\text{ref}}$  and the best-performing case  $E_c^{\text{opt}}$  are marked with a blue cross and a red star, respectively. The energy spectra are calculated from the whole temporal profile depicted in **a**).

In the first iteration, all batch sizes identified the same best-performing point, providing an improvement of more than 30× compared to the reference. Since only one simulation was required,  $N = 1$  was the most efficient at this stage in terms of computational resources. However, after the next seven iterations, the  $N = 1$  case reached values comparable to the second-iteration results of  $N = 4$  and  $N = 8$ , which means that for a similar number of simulations, all batch sizes gave similar results.

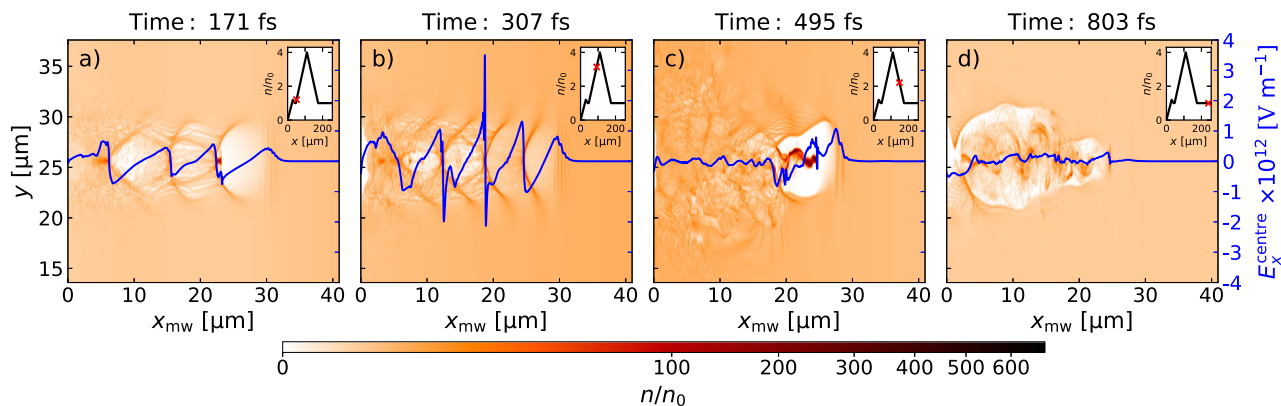
Figure 2b shows the wall-clock time measured for each batch, where the time corresponds to the interval between the start of the first simulation in the batch and the completion of the longest-running one. Overall, the total times are comparable because the simulations were run in parallel. The average wall-clock time per iteration is 4.4, 4.1, and 5.5 h for  $N = 1, 4$ , and 8, respectively.

Cases with better-performing cost values tend to run longer, as they typically involve processing more electrons, as discussed later. In general, however, most variations in wall-clock time across iterations and batch sizes were driven by fluctuating queue delays on the supercomputing system rather than by differences among the simulations themselves. This shows that the performance of BBO depends on available parallelization. Note that increasing the batch size leads to higher total CPU usage, since more simulations run simultaneously.

Across all batch sizes  $N = 1, 4$ , and 8, the best-performing parameters varied only slightly,  $d_u = 6.8 - 8.4 \mu\text{m}$ ,  $d_s = 112 - 127 \mu\text{m}$ , while  $n_p$  remained consistently at  $4n_0$ . We performed a local parameter scan around the best-performing optimized point across all batch sizes, for which  $d_u = d_u^{\text{opt}} = 6.8 \mu\text{m}$  and  $d_s = d_s^{\text{opt}} = 127.4 \mu\text{m}$  and  $n_p = n_p^{\text{opt}} = 4.0 n_0$ , found in the

7th iteration of a batch search with  $N = 8$ . The scan, depicted in Fig. 3 shows that  $d_u$  is a highly sensitive parameter. Once it exceeds roughly the bubble-scale length at  $n_0$  ( $\lambda_p = 7.5 \mu\text{m}$ ), the cost drops rapidly below  $20 C^{\text{ref}}$  (Fig. 3a), which shows that the spike must follow the initial gradient within approximately one bubble scale length. In contrast,  $d_s$  and  $n_p$  (Fig. 3b, c) show a much more gradual decrease and remain robust across a wider range. Aside from slightly underestimating the performance at larger  $d_s$ , overall, the final BBO model agrees well with the evaluated data. Note that, as shown in Supplementary Note 1, Fig. S1, the uncertainty remains large outside this vicinity, around  $\sigma_{\hat{C}} \approx 14$ . Therefore, it is generally possible that additional enhancement regions could have been found outside the best-performing region identified by our model if the exploration had been boosted by slightly modifying the optimization strategy. Exploring such possibilities lies beyond the scope of the present work.

The characteristics of the radiation for two simulations, the best-performing optimized point  $\{d_u^{\text{opt}}, d_s^{\text{opt}}, n_p^{\text{opt}}\}$  and the reference case without the spike are shown in Fig. 4. Figure 4a depicts the temporal profile of the radiation and insets Fig. 4b and c show corresponding zoomed-in profiles around the peak values for the two studied cases. As can be seen in Fig. 4a, the optimized pulse arrives earlier in time compared to the reference case, and its peak value is higher by more than 25×. The value of  $W_{50}$  increased to  $W_{50}^{\text{opt}} = 305.0 \text{ nJ sr}^{-1}$  and the value of  $\tau_{50}$  dropped to  $\tau_{50}^{\text{opt}} = 43.50 \text{ as}$ . The optimized case also exhibits a more pronounced tail-like structure, with oscillations of a magnitude comparable to the reference pulse itself. As a consequence, the total radiated energy naturally increases, as shown in Fig. 4d. Importantly, not only does the low-energy part of the spectrum



**Fig. 5 | Plasma density evolution for the best-performing case.** **a** Early stage of propagation shortly after the first electron injection. **b** Propagation through an increasing gradient of the density spike. **c** A second injection is triggered while the laser pulse propagates through the decreasing gradient. **d** End of the acceleration and degradation of the bubble structure. The density profiles are shown at the center of the  $z$ -axis,  $z = 25.6 \mu\text{m}$ , and  $x_{\text{mw}}$  is the  $x$ -coordinate comoving with the simulation window. For better visualization,  $n/n_0$  is transformed using a power-law function

with exponent 0.35,  $((n/n_0)/(n/n_0)_{\text{max}})^{0.35}$ , where  $(n/n_0)_{\text{max}}$  corresponds to the maximum value of the normalized density  $n/n_0$ .  $E_x^{\text{centre}}$  (blue lines) corresponds to the longitudinal electric field at the center of the  $y$  and  $z$ -axis,  $y = z = 25.6 \mu\text{m}$ . The red crosses in the insets show the corresponding pulse position within the density gradient. The times shown above the panels are measured from when the pulse center reaches the simulation box entrance.

increase, but higher energies compared to the reference were also reached. The critical energy increases from  $E_c^{\text{ref}} = 0.26 \text{ keV}$  for the reference case to  $E_c^{\text{opt}} = 0.60 \text{ keV}$  for the best-performing case.

### Analysis of the betatron enhancement

To obtain insight into the mechanism underlying the improvement of the betatron radiation, we analyze the dynamics of the electron beam propagating through the density spike. The evolution of the plasma density profile for the best-performing case is shown in Fig. 5. Shortly after the first plasma gradient that induces the electron injection, a short electron beam is observed at the rear part of the bubble (Fig. 5a). As the laser pulse propagates further through the increasing density of the spike, the bubble shrinks and electrons from the first injection are no longer sustained in the first bubble (Fig. 5b). After the peak of the spike, another electron injection with an abundance of electrons is observed (Fig. 5c). This injection leads to an intense attosecond betatron radiation burst with a noisy radiation tail (seen in Fig. 4a and b). At the end of the acceleration, the original bubble structure is destroyed (Fig. 5d). The leading part of the novel electron beam generates its own plasma wakefield, with a stronger electric field than the one of the laser pulse. The suppression of the laser wakefield at the center of the  $y$  and  $z$  axes,  $E_x^{\text{centre}}$ , can be seen in Fig. 5d, when compared to Fig. 5a–c, where the wakefield originating from the laser pulse location is visible. The electron beam tail is accelerated in the plasma wakefield created by the leading electron beam front. A small overlap of the electron beam and laser tail was present during the final tens of femtoseconds and might have contributed to the additional radiation gain. More details on the electron and betatron beam properties, including electron energy, charge, the angular radiation spectrum, as well as the effects of multiple wakefield periods and a plasma density down-ramp at the plasma exit, are described in Supplementary Note 2.

### Discussion

Batch Bayesian optimization of attosecond betatron pulses, covering a multi-dimensional optimization space parametrizing a plasma source with a density spike, identified a parameter region with significant radiation enhancement. Comparing the performance for the batch sizes of  $N = 1, 4$ , and  $8$ , we observed that batch evaluations with  $N > 1$  can substantially accelerate the optimization relative to the sequential approach ( $N = 1$ ), assuming that the system supports high parallelization. However, the total available resources must be carefully considered, as larger batch sizes require more simulations per iteration. In our case,  $N = 4$  offers the best compromise between wall-clock time, simulation count, and achieved improvement. The

optimal choice of batch size, however, is inherently problem-dependent and varies depending on the computational platform and overall optimization goals.

BBO found a region in the parameter space where the cost function increased more than 60 times compared to the reference value without the density spike. This region corresponds to a long, high-density spike positioned shortly after the initial gradient. The analysis of the PIC simulations revealed that the enhancement was due to a second, additional electron injection process induced by the density spike.

These findings highlight the efficiency of the Bayesian approach for multiparametric scans in optimization problems, especially when the number of simulations is limited or when the output is complex. In particular, revealing the efficiency of high-density down-ramp electron injection in improving the betatron pulse characteristics provides valuable guidance for experimental investigation. It should be emphasized that although the optimization converged within the explored domain, global optimality cannot be proven in principle for such multi-dimensional problems that may contain several local extrema. This limitation is inherent to such complex optimization landscapes. Choosing the batch approach instead of sequential, in particular, is well motivated if queuing and running the simulations is time-consuming, or when time-intensive manual processing is required. Future work could also benefit from using multi-objective Bayesian optimization techniques to better capture trade-offs across the explored parameter space.

### Methods

#### Particle-in-cell simulations

The setup for the Smilei PIC simulations was as follows: The simulation box has dimensions  $L_x \times L_y \times L_z = 40.96 \mu\text{m} \times 51.2 \mu\text{m} \times 51.2 \mu\text{m}$ . A moving-window technique was applied, with a velocity of  $c$  in the  $x$  direction. The plasma was assumed to be fully ionized, with only electrons initialized in the simulation. The ions were considered immobile and were thus not initialized, in order to reduce computational time. The simulation grid was discretized with a resolution of  $\Delta x \times \Delta y \times \Delta z = 17.8 \text{ nm} \times 133.3 \text{ nm} \times 133.3 \text{ nm}$ , and the simulation timestep was  $\Delta t_{\text{SMILEI}} = 57 \text{ as}$ . One electron macroparticle per cell was used, with regular position initialization. The Maxwell equations were solved using the Lehe scheme<sup>59</sup> in combination with a binomial filter<sup>60</sup>. One macroparticle represents a set of electrons with a corresponding statistical weight. The numerical convergence tests, showing that this approach is sufficient to capture the physics for the optimization, are described in more detail in Supplementary Note 3.

### Betatron radiation calculation

Betatron radiation was computed using the FIKA code, which evaluates far-field emission via the Liénard-Wiechert potentials, using an approach related to previous work<sup>61</sup>. FIKA requires the trajectories and momenta of the particles as input. In FIKA, only the incoherent part of the radiation is calculated, neglecting any coherence effects in the radiation. This approach is sufficient, as electrons emit radiation incoherently as betatron radiation from relativistic electron beams<sup>27</sup>, and collective coherence effects can be neglected. Each PIC macroparticle is treated as a statistically independent radiation source, with no interference between different macroparticles. The final radiation is then obtained through a weighted summation of the radiation contributions from each PIC macroparticle<sup>62</sup>, with macroparticle statistical weights incorporated into the calculation.

Particle trajectories were extracted from the output of the Smilei code, sampled every  $15\Delta t_{\text{SMILEI}}$ , where electron positions and momenta were tracked for each electron macroparticle with energy greater than 10 MeV. FIKA has an additional module that converts a Smilei output file containing particle trajectories into a format compatible with FIKA, which is also available in the FIKA GitHub repository<sup>58</sup>. The code first computed  $\frac{d^2W}{d\Omega dt} \Big|_{\text{on-axis}}$ , evaluated at an observer position of  $(x, y, z) = (1 \text{ m}, 0, 0)$ . To obtain  $\frac{d^2W}{d\Omega dt} \Big|_{\text{on-axis}}$ , the temporal resolution was refined from  $15\Delta t_{\text{SMILEI}}$  to  $\Delta t_{\text{FIKA}} = 15\Delta t_{\text{SMILEI}} / (2\gamma_{e\text{max}}^2)$ , where  $\gamma_{e\text{max}}$  is the maximum electron Lorentz factor along each trajectory. The radiation signal was then interpolated onto this finer temporal grid before performing the Fourier transform.

### Batch Bayesian Optimization

Multi-objective Bayesian optimization is guided by an empirically defined cost function that quantifies performance across the parameter space. A Gaussian Process (GP) was used as a surrogate model to approximate the cost function. In a GP model, correlations between input points are specified by a kernel function. Here, we used the Matérn 5/2 kernel, (defined in Supplementary Note 1), which corresponds to the prior assumption that the cost function is twice mean-square differentiable. Matérn kernels enable more relaxed assumptions on the smoothness of the cost function, which allows to capture also sharper variations that might locally occur, compared to, for instance, the squared exponential (RBF) kernel, which assumes infinitely smooth functions. The acquisition function guides the search by balancing exploration and exploitation to determine where the next point should be evaluated. Here, we used the Expected Improvement (EI) acquisition function<sup>63,64</sup>, which estimates the expected gain over the current best observed value. Batch points were selected using the Fantasizer approach<sup>65</sup>, which sequentially maximizes EI while updating the GP with “fantasy” observations to avoid selecting points that are too close together.

The optimization was implemented using the Trieste library<sup>66</sup> and its Efficient Global Optimization framework. The Bayesian optimization process was performed as follows. (i) A batch of  $N$  PIC simulations was performed with the Smilei code using selected values for  $\{d_w, d_s, n_p\}$ . Using the Smilei output of particle trajectories,  $W_{50}$  and  $\tau_{50}$  were obtained with the FIKA code. (ii) The cost function values were then computed for each point according to Eq. (1). Based on the new results, we updated the surrogate model of the cost function with GP. (iii) The acquisition function was then used to propose  $N$  new points. Candidate batches were checked against the constraints in Eqs. (2)–(5), and if fewer than  $N$  valid points were found, additional candidates were generated until a full batch of  $N$  valid points was obtained. (iv) Steps (i)–(iii) were repeated for 8 iterations.

We assumed noise-free simulation outputs from Smilei and FIKA. In the PIC Smilei simulations, macroparticles were initialized with regular positions and cold momentum, avoiding randomized initialization. We neglected potential statistical variability between runs, as all simulations were performed with the same number of particles per cell, identical resolution, and on the same computational system (16 nodes, 128 cores each), ensuring reproducible outputs, with any remaining variability expected to be negligible. The GP surrogate model was therefore trained assuming a

noise-free problem, with a very small likelihood (noise) variance of  $10^{-7}$  included for numerical stability.

### Data availability

The data that support the findings of this study are available from the corresponding author upon request.

### Code availability

The FIKA code developed for this work is available on GitHub under the MIT License<sup>58</sup>.

Received: 8 May 2025; Accepted: 3 February 2026;

Published online: 18 February 2026

### References

- McPherson, A. et al. Studies of multiphoton production of vacuum-ultraviolet radiation in the rare gases. *J. Optical Soc. Am. B* **4**, 595–601 (1987).
- Ferray, M. et al. Multiple-harmonic conversion of 1064 nm radiation in rare gases. *J. Phys. B: At. Mol. Optical Phys.* **21**, L31 (1988).
- L’Huillier, A., Schafer, K. & Kulander, K. Higher-order harmonic generation in xenon at 1064 nm: the role of phase matching. *Phys. Rev. Lett.* **66**, 2200 (1991).
- Lewenstein, M., Balcou, P., Ivanov, M. Y., L’Huillier, A. & Corkum, P. B. Theory of high-harmonic generation by low-frequency laser fields. *Phys. Rev. A* **49**, 2117 (1994).
- Horný, V., Krus, M., Yan, W. & Fülöp, T. Attosecond betatron radiation pulse train. *Sci. Rep.* **10**, 15074 (2020).
- Ferri, J., Horný, V. & Fülöp, T. Generation of attosecond electron bunches and x-ray pulses from few-cycle femtosecond laser pulses. *Plasma Phys. Control. Fusion* **63**, 045019 (2021).
- Tajima, T. & Dawson, J. M. Laser electron accelerator. *Phys. Rev. Lett.* **43**, 267 (1979).
- Luttikhof, M., Khachatryan, A., Van Goor, F. & Boller, K.-J. Generating ultrarelativistic attosecond electron bunches with laser wakefield accelerators. *Phys. Rev. Lett.* **105**, 124801 (2010).
- Tooley, M. et al. Towards attosecond high-energy electron bunches: controlling self-injection in laser-wakefield accelerators through plasma-density modulation. *Phys. Rev. Lett.* **119**, 044801 (2017).
- Zhao, Q. et al. Sub-femtosecond electron bunches in laser wakefield acceleration via injection suppression with a magnetic field. *Plasma Phys. Control. Fusion* **61**, 085015 (2019).
- Kim, J., Wang, T., Khudik, V. & Shvets, G. Subfemtosecond wakefield injector and accelerator based on an undulating plasma bubble controlled by a laser phase. *Phys. Rev. Lett.* **127**, 164801 (2021).
- Deng, A., Li, X., Luo, Z., Li, Y. & Zeng, J. Generation of attosecond micro bunched beam using ionization injection in laser wakefield acceleration. *Opt. Express* **31**, 19958–19967 (2023).
- Tomassini, P., Horny, V. & Doria, D. Attosecond pulses from ionization injection wakefield accelerators. *Instruments* **7**, 34 (2023).
- Sun, T., Zhao, Q., Wan, F., Salamin, Y. I. & Li, J.-X. Generation of ultrabright polarized attosecond electron bunches via dual-wake injection. *Phys. Rev. Lett.* **132**, 045001 (2024).
- Tomassini, P. et al. Ultra-high-brightness and tuneable attosecond-long electron beams with the laser wake field acceleration. *Sci. Rep.* **15**, 40794 (2025).
- Albert, F. & Thomas, A. G. Applications of laser wakefield accelerator-based light sources. *Plasma Phys. Control. Fusion* **58**, 103001 (2016).
- Nemeth, K. et al. Laser-driven coherent betatron oscillation in a laser-wakefield cavity. *Phys. Rev. Lett.* **100**, 095002 (2008).
- Cipiccia, S. et al. Gamma-rays from harmonically resonant betatron oscillations in a plasma wake. *Nat. Phys.* **7**, 867–871 (2011).
- Kneip, S. et al. X-ray phase contrast imaging of biological specimens with femtosecond pulses of betatron radiation from a compact laser plasma wakefield accelerator. *Appl. Phys. Lett.* **99**, 093701 (2011).

20. Fourmaux, S., Hallin, E., Chaulagain, U., Weber, S. & Kieffer, J. Laser-based synchrotron X-ray radiation experimental scaling. *Opt. Express* **28**, 3147–3158 (2020).
21. Cole, J. et al. Laser-wakefield accelerators as hard X-ray sources for 3D medical imaging of human bone. *Sci. Rep.* **5**, 13244 (2015).
22. Hussein, A. E. et al. Laser-wakefield accelerators for high-resolution X-ray imaging of complex microstructures. *Sci. Rep.* **9**, 3249 (2019).
23. Seres, E., Seres, J. & Spielmann, C. X-ray absorption spectroscopy in the keV range with laser-generated high harmonic radiation. *Appl. Phys. Lett.* **89**, 181919 (2006).
24. Yamada, S. et al. Broadband high-energy resolution hard X-ray spectroscopy using transition edge sensors at Spring-8. *Rev. Sci. Instrum.* **92**, 013103 (2021).
25. Kiselev, S., Pukhov, A. & Kostyukov, I. X-ray generation in strongly nonlinear plasma waves. *Phys. Rev. Lett.* **93**, 135004 (2004).
26. Rousse, A. et al. Production of a keV X-Ray beam from synchrotron radiation in relativistic laser-plasma interaction. *Phys. Rev. Lett.* **93**, 135005 (2004).
27. Corde, S. et al. Femtosecond X rays from laser-plasma accelerators. *Rev. Mod. Phys.* **85**, 1–48 (2013).
28. Huang, K. et al. Resonantly enhanced betatron hard x-rays from ionization injected electrons in a laser plasma accelerator. *Sci. Rep.* **6**, 27633 (2016).
29. Mangles, S. P. et al. Laser-wakefield acceleration of monoenergetic electron beams in the first plasma-wave period. *Phys. Rev. Lett.* **96**, 215001 (2006).
30. Chen, L. et al. Bright betatron X-ray radiation from a laser-driven-clustering gas target. *Sci. Rep.* **3**, 1912 (2013).
31. Mangles, S. et al. Controlling the spectrum of x-rays generated in a laser-plasma accelerator by tailoring the laser wavefront. *Appl. Phys. Lett.* **95**, <https://pubs.aip.org/aip/apl/article/95/18/181106/131664/Controlling-the-spectrum-of-x-rays-generated-in-a> (2009).
32. Popp, A. et al. All-optical steering of laser-wakefield-accelerated electron beams. *Phys. Rev. Lett.* **105**, 215001 (2010).
33. Wood, J. et al. Enhanced betatron radiation from a laser wakefield accelerator in a long focal length geometry. *Gas* **1**, 2 (2017).
34. Yan, W. et al. Concurrence of monoenergetic electron beams and bright X-rays from an evolving laser-plasma bubble. *Proc. Natl. Acad. Sci.* **111**, 5825–5830 (2014).
35. Zhang, Z. et al. Enhanced x-rays from resonant betatron oscillations in laser wakefield with external wigglers. *Plasma Phys. Control. Fusion* **58**, 105009 (2016).
36. Ho, Y.-C. et al. Induction of electron injection and betatron oscillation in a plasma-waveguide-based laser wakefield accelerator by modification of waveguide structure. *Phys. Plasmas* **20**, 083104 (2013).
37. Wallin, E., Gonoskov, A. & Marklund, M. Radiation emission from braided electrons in interacting wakefields. *Phys. Plasmas* **24**, <https://doi.org/10.1063/1.4997440> (2017).
38. Shaw, J. et al. Role of direct laser acceleration in energy gained by electrons in a laser wakefield accelerator with ionization injection. *Plasma Phys. Control. Fusion* **56**, 084006 (2014).
39. Zhao, T. et al. High-flux femtosecond x-ray emission from controlled generation of annular electron beams in a laser wakefield accelerator. *Phys. Rev. Lett.* **117**, 094801 (2016).
40. Ferri, J. & Davoine, X. Enhancement of betatron X rays through asymmetric laser wakefield generated in transverse density gradients. *Phys. Rev. Accel. Beams* **21**, 091302 (2018).
41. Ma, Y. et al. Angular streaking of betatron X-rays in a transverse density gradient laser-wakefield accelerator. *Phys. Plasmas* **25**, <https://pubs.aip.org/aip/pop/article/25/11/113105/263228/Angular-streaking-of-betatron-X-rays-in-a> (2018).
42. Ta Phuoc, K. et al. Betatron radiation from density-tailored plasmas. *Phys. Plasmas* **15**, <https://pubs.aip.org/aip/pop/article/15/6/063102/923579/Betatron-radiation-from-density-tailored-plasmas> (2008).
43. Guo, B. et al. Enhancement of laser-driven betatron x-rays by a density-depressed plasma structure. *Plasma Phys. Control. Fusion* **61**, 035003 (2019).
44. Ferri, J. et al. High-brilliance betatron  $\gamma$ -ray source powered by laser-accelerated electrons. *Phys. Rev. Lett.* **120**, 254802 (2018).
45. Döpp, A. et al. Energy boost in laser wakefield accelerators using sharp density transitions. *Phys. Plasmas* **23**, <https://pubs.aip.org/aip/pop/article/23/5/056702/966617/Energy-boost-in-laser-wakefield-accelerators-using> (2016).
46. Shalloo, R. et al. Automation and control of laser wakefield accelerators using Bayesian optimization. *Nat. Commun.* **11**, 6355 (2020).
47. Jalas, S. et al. Bayesian optimization of a laser-plasma accelerator. *Phys. Rev. Lett.* **126**, 104801 (2021).
48. Jalas, S. et al. Tuning curves for a laser-plasma accelerator. *Phys. Rev. Accel. Beams* **26**, 071302 (2023).
49. Irshad, F. et al. Pareto optimization and tuning of a laser wakefield accelerator. *Phys. Rev. Lett.* **133**, 085001 (2024).
50. Ferran Pousa, A. et al. Bayesian optimization of laser-plasma accelerators assisted by reduced physical models. *Phys. Rev. Accel. Beams* **26**, 084601 (2023).
51. Irshad, F., Karsch, S. & Döpp, A. Multi-objective and multi-fidelity Bayesian optimization of laser-plasma acceleration. *Phys. Rev. Res.* **5**, 013063 (2023).
52. Zhong, J. et al. Simulation of laser plasma wakefield acceleration with external injection based on Bayesian optimization. *Plasma Sci. Technol.* <http://iopscience.iop.org/article/10.1088/2058-6272/ad91e8> (2024).
53. Valenta, P., Esirkepov, T. Z., Ludwig, J. D., Wilks, S. C. & Bulanov, S. V. Bayesian optimization of electron energy from laser wakefield accelerator. *Phys. Rev. Accel. Beams* **28**, 094601 (2025).
54. Ye, H. et al. Fast optimization for betatron radiation from laser wakefield acceleration based on Bayesian optimization. *Results Phys.* **43**, 106116 (2022).
55. Bulanov, S., Naumova, N., Pegoraro, F. & Sakai, J. Particle injection into the wave acceleration phase due to nonlinear wake wave breaking. *Phys. Rev. E* **58**, R5257 (1998).
56. Lu, W. et al. Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime. *Phys. Rev. Spec. Top. Accel. Beams* **10**, 061301 (2007).
57. Derouillat, J. et al. Smilei: A collaborative, open-source, multi-purpose particle-in-cell code for plasma simulation. *Comput. Phys. Commun.* **222**, 351–373 (2018).
58. Maslarova, D. & Hansson, A. chalmersplasmtheory/fika: Version 1.0 <https://doi.org/10.5281/zenodo.15350881> (2025).
59. Lehe, R., Lifschitz, A., Thaur, C., Malka, V. & Davoine, X. Numerical growth of emittance in simulations of laser-wakefield acceleration. *Phys. Rev. ST Accel. Beams* **16**, 021301 (2013).
60. Vay, J.-L., Geddes, C., Cormier-Michel, E. & Grote, D. Numerical methods for instability mitigation in the modeling of laser wakefield accelerators in a Lorentz-boosted frame. *J. Comput. Phys.* **230**, 5908–5929 (2011).
61. Horný, V. et al. Temporal profile of betatron radiation from laser-driven electron accelerators. *Phys. Plasmas* **24**, 063107 (2017).
62. Pausch, R. et al. Quantitatively consistent computation of coherent and incoherent radiation in particle-in-cell codes—a general form factor formalism for macro-particles. *Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip.* **909**, 419–422 (2018).
63. Močkus, J. On Bayesian methods for seeking the extremum. In *Proc. Optimization Techniques IFIP Technical Conference Novosibirsk*, July 1–7, 1974, Vol. 6, 400–404, [https://link.springer.com/chapter/10.1007/3-540-07165-2\\_55](https://link.springer.com/chapter/10.1007/3-540-07165-2_55) (Springer, 1975).
64. Jones, D. R., Schonlau, M. & Welch, W. J. Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13**, 455–492 (1998).

65. Ginsbourger, D., Le Riche, R. & Carraro, L. Kriging is well-suited to parallelize optimization. In *Computational Intelligence in Expensive Optimization Problems*, 131–162, [https://link.springer.com/chapter/10.1007/978-3-642-10701-6\\_6](https://link.springer.com/chapter/10.1007/978-3-642-10701-6_6) (Springer, 2010).
66. SecondMind Labs. Batch Bayesian optimization. [https://secondmind-labs.github.io/trieste/4.2.2/notebooks/batch\\_optimization.html](https://secondmind-labs.github.io/trieste/4.2.2/notebooks/batch_optimization.html) (2024).

## Acknowledgements

The authors are grateful to Patrik Jansson and Ida Ekmark from Chalmers University of Technology, Miroslav Krus from the Institute of Plasma Physics of the Czech Academy of Sciences, Sarah Newton from UKAEA and David Gregocki from CNR—Istituto Nazionale di Ottica for fruitful discussions. This project received funding from the Knut and Alice Wallenberg Foundation (Grants Nos. KAW 2020.0111 and 2023.0249). The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement No. 2022-06725, and by EuroHPC Joint Undertaking through access to Karolina at IT4Innovations (VŠB-TU), Czechia under project number EHP-REG-2025R01-007, together with Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90140). V.H. draws support from the European Union, the Romanian Government and the Health Program, within the project SMIS Code: 326475, and the Romanian Ministry of Research, Innovation and Digitalization: Program Nucleu PN23210105.

## Author contributions

D.M., J.F. and I.P. conceived the main idea with inputs from V.H. and M.L. D.M., A.H., and M.L. designed the numerical simulations. D.M. and A.H. conducted the numerical simulations and analyzed the results. D.M., T.F. and I.P. wrote the manuscript with inputs from V.H. and M.L.

## Funding

Open access funding provided by Chalmers University of Technology.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s42005-026-02542-6>.

**Correspondence** and requests for materials should be addressed to Dominika Maslarova.

**Peer review information** *Communications Physics* thanks Davorin Peceli and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2026