



PAPER

Multi-parameter Bayesian optimisation of laser-driven ion acceleration in particle-in-cell simulations

OPEN ACCESS

RECEIVED
4 April 2022REVISED
14 June 2022ACCEPTED FOR PUBLICATION
1 July 2022PUBLISHED
19 July 2022

Original content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution
of this work must
maintain attribution to
the author(s) and the
title of the work, journal
citation and DOI.

E J Dolier¹ , M King^{1,2}, R Wilson¹, R J Gray¹  and P McKenna^{1,2,*} ¹ SUPA Department of Physics, University of Strathclyde, Glasgow G4 0NG, United Kingdom² The Cockcroft Institute, Sci-Tech Daresbury, Warrington WA4 4AD, United Kingdom

* Author to whom any correspondence should be addressed.

E-mail: paul.mckenna@strath.ac.uk**Keywords:** laser-driven ion acceleration, Bayesian optimization, machine learningSupplementary material for this article is available [online](#)

Abstract

High power laser-driven ion acceleration produces bright beams of energetic ions that have the potential to be applied in a wide range of sectors. The routine generation of optimised and stable ion beam properties is a key challenge for the exploitation of these novel sources. We demonstrate the optimisation of laser-driven proton acceleration in a programme of particle-in-cell simulations controlled by a Bayesian algorithm. Optimal laser and plasma conditions are identified four times faster for two input parameters, and approximately one thousand times faster for four input parameters, when compared to systematic, linear parametric variation. In addition, a non-trivial optimal condition for the front surface density scale length is discovered, which would have been difficult to identify by single variable scans. This approach enables rapid identification of optimal laser and target parameters in simulations, for use in guiding experiments, and has the potential to significantly accelerate the development and application of laser-plasma-based ion sources.

1. Introduction

The interaction of high power laser pulses with plasma results in the generation of energetic pulses of ions with unique properties, such as high brightness and laminarity [1, 2]. Due to the high electric fields achieved in plasma, these novel ion sources are inherently compact and could bring new capabilities to accelerator science, with wide-ranging applications in areas as diverse as medicine [3–7], industry [8–10], and inertial fusion energy [11, 12]. Optimising, stabilising and controlling the unique properties of the ion beams is an outstanding challenge that needs to be addressed for the transformative potential of these novel sources to be realised.

Of the laser-driven ion acceleration mechanisms explored to date, target normal sheath acceleration (TNSA) [13, 14] is the most widely investigated, and is readily accessible with current high-power laser technology. In this scheme, which typically involves a thin foil target, a relativistically intense ($>10^{18}$ W cm⁻² for laser wavelength of the order of 1 μ m) laser pulse irradiates the target, producing a dense plasma and transferring energy to plasma electrons that propagate within the target (so-called fast electrons) and generate TV m⁻¹ electrostatic sheath fields at the target–vacuum interfaces. The field at the target rear side accelerates ions (predominately protons due to their high charge-to-mass ratio) in the forward direction, to form a beam with a broad energy distribution [15]. The maximum ion energy depends on parameters such as the temperature of the fast electron population [16, 17], which is strongly dependent on the laser irradiance, $I_L \lambda_L^2$, where I_L and λ_L are the laser intensity and wavelength, respectively [18–21]. Increasing the laser intensity and decreasing the target thickness or density enables new ion acceleration mechanisms to become accessible. At sufficiently high intensity, the radiation pressure exerted by the laser light incident on the plasma exceeds its thermal pressure, resulting in the laser boring into the plasma [18]. The resulting displacement of electrons establishes a strong electric field that accelerates ions

via so-called radiation pressure acceleration (RPA) [22–24]. This approach can, in principle, produce a peaked ion energy distribution, with a much faster maximum energy scaling with intensity, and a high laser-to-ion energy conversion efficiency, compared to TNSA [25]. The light-sail mode of RPA, which occurs in ultrathin foils, is particularly promising in this respect [22, 26, 27]. High proton energies are produced in ultrathin foils that become relativistically transparent to the laser light during the interaction, as this enables additional heating of the plasma electrons, enhancing the electric fields and thus ion energies. There are several variations of this scheme, including the laser break-out afterburner [28] mechanism, the transparency-enhanced hybrid RPA–TNSA scheme [29] and the synchronised acceleration by slow light [30].

Common to all of these laser-driven ion acceleration mechanisms is the need to simultaneously optimise multiple laser and plasma parameters to tailor the properties of the resultant ion beam. Progress to date has involved the sequential scanning of a selected laser or target parameter while others are fixed at expected optimal values. The recent development of high-repetition-rate (multi-Hz), high-intensity, short-pulse lasers based on Ti:sapphire technology will accelerate the development of laser-driven ion sources by generating experimental data at a rate that facilitates live statistical analysis and feedback [31]. However, some applications (e.g. inertial fusion studies and some topics in high energy density science) require high energy (Nd:glass-based) lasers or specialist dense targets, which inherently limit repetition rates and restrict the application of optimisation schemes based solely on statistical analysis of real-time experimental data. These sources can benefit from rapid developments in high performance computing (HPC) that are enabling much higher volumes of complementary simulation data, greatly increasing the parameter space that can be explored via modelling. These parallel achievements have the potential to transform our approach to developing laser-driven ion acceleration by enabling the application of machine learning (ML) techniques to address this challenge.

The ML technique of Bayesian optimisation [32] is particularly useful when limited data is available and where the target function to be optimised returns noisy evaluations. This is the case in experiments with low repetition rate lasers, where the overall shot numbers are small and may be subject to shot-to-shot fluctuations. It is also true in numerical investigations involving high-fidelity particle-in-cell (PIC) simulations, for which the number of simulations that can be run is limited by the large computational and storage resources required. By contrast, ML approaches based on neural networks and genetic algorithms require large quantities of data to make accurate predictions, and direct optimisation algorithms can struggle to optimise noisy functions. Bayesian optimisation has been applied to laser-driven electron acceleration experiments [33–35], and Bayesian statistics have been used to model laser-driven ion acceleration studies [36], along with other common ML techniques [37, 38].

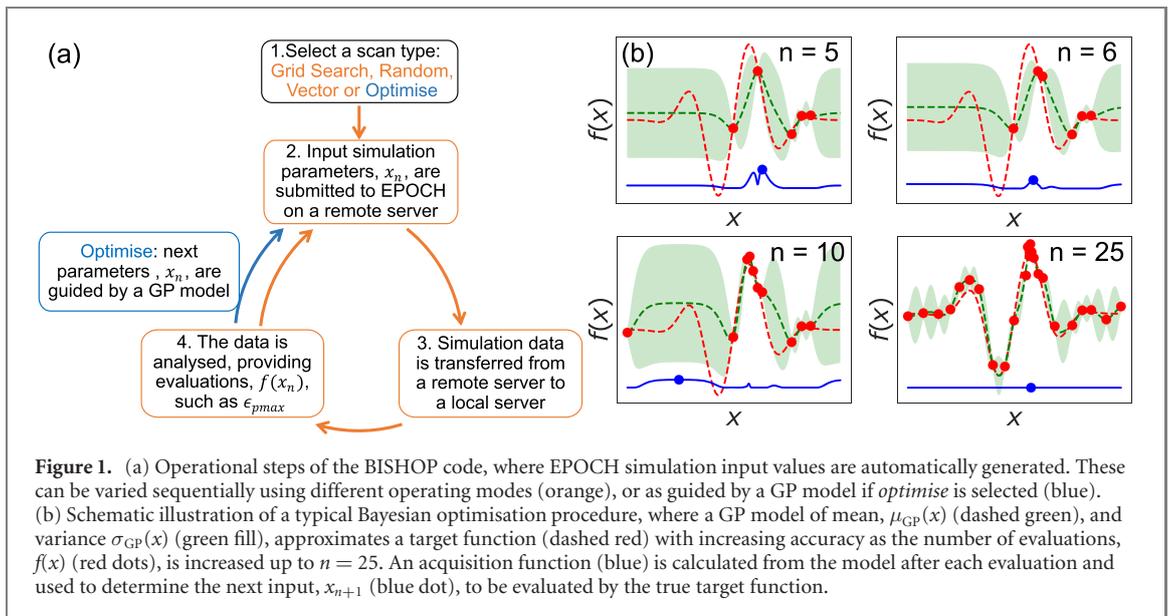
In this paper, we demonstrate the application of a Bayesian optimisation algorithm to guide the choice of input conditions and analyse output data from simulations performed using a PIC code, enabling the maximum energy to which protons are accelerated by TNSA to be optimised as a function of four key laser and target parameters. This was achieved in a fraction of the time and the number of simulations required compared to an approach involving systematic, linear parametric variation. The results demonstrate the potential to apply Bayesian algorithm techniques to laser-driven ion acceleration in situations where limited data is available and where multi-parameter optimisation would not otherwise be possible. We also report the development of a code called BISHOP, which facilitated this work by automatically generating simulation inputs, after analysing output data from previous simulations, and submitting them to a remote HPC cluster. Using the BISHOP code significantly reduces the time taken to produce and prepare simulation data for use in a ML model.

2. Methodology

2.1. Automated linear scanning of simulation input parameters

Before considering Bayesian optimisation, a reference data-set of one hundred 2D PIC simulations was produced, in which input parameters, x , were linearly varied in a conventional *grid search* operation (systematic parameter scans in fixed increments), and submitted to the fully relativistic EPOCH PIC code [39]. This process was automated using the BISHOP code, which generates the input files for EPOCH and submits them to a remote HPC cluster, before retrieving and evaluating the output $f(x_n)$ for each increment n , before updating the inputs. In *grid search* mode, $f(x_n)$ is a defined output, e.g. maximum proton energy (ϵ_{pmax}). The operational steps in this mode are depicted schematically (in orange) in figure 1(a).

To begin this process, the user enters a base set of simulation input parameters, together with minimum and maximum bounds for the key parameters to be varied. Each simulation in the reference data set was initialised with an average of ten particles per cell, per species, and grid dimensions of $16\,000 \times 2400$ cells

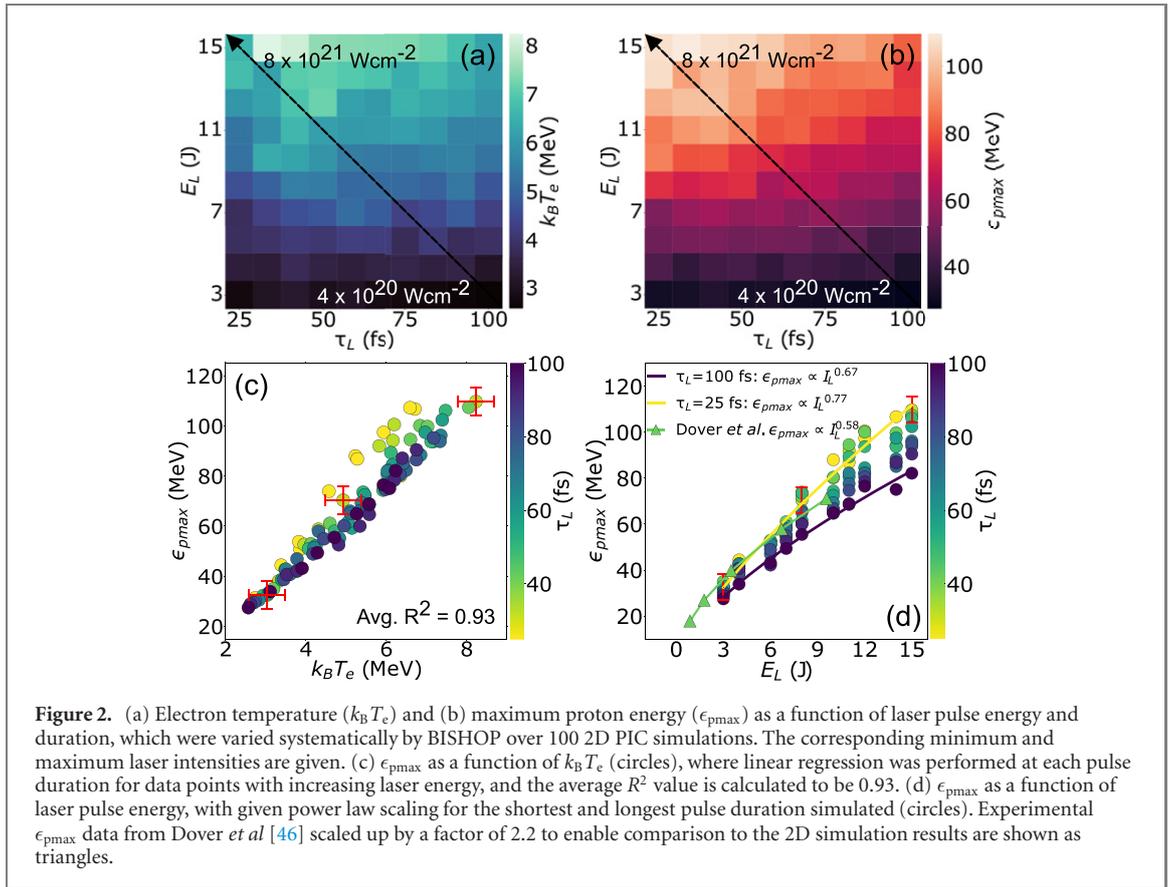


were defined, corresponding to a domain of $(-10 \rightarrow +50) \mu\text{m}$ and $(-20 \rightarrow +20) \mu\text{m}$ in the X and Y dimensions, respectively. The target thickness was defined by a $1 \mu\text{m}$ region in X containing H^+ and C^{6+} ions with a starting temperature of 0.04 keV and electrons with a starting temperature of 10 keV . The initial electron density was selected as $203n_{\text{crit}}$, where n_{crit} is the plasma critical density, defined as $n_{\text{crit}} = \frac{m_e \epsilon_0 \omega_L^2}{e}$, and m_e , ϵ_0 , ω_L , and e are the electron mass, free-space permittivity, laser angular frequency, and the elementary charge, respectively. This initial electron density was chosen to approximate a solid plastic target by neutralising an equal mix of H^+ and C^{6+} ions. The laser energy, E_L , and pulse duration, τ_L , were selected as key parameters, with bounds of $(0.3 \rightarrow 15.0) \text{ J}$ and $(25 \rightarrow 100) \text{ fs}$, respectively. This reflects typical conditions available with Ti:sapphire high power lasers at present, as well as those coming online in the near future [40–42]. For these ranges of parameters, TNSA is the dominant ion acceleration mechanism and is responsible to the highest energy protons. Example simulation results verifying this are provided in the supplementary material (<https://stacks.iop.org/NJP/24/073025/mmedia>).

Over 100 simulations, the laser energy and pulse duration were varied linearly in equally spaced values between these bounds, and the focal spot size was fixed at $\phi_L = 3 \mu\text{m}$. Upon completion of each simulation, the output data was transferred by the BISHOP code from the remote HPC cluster to a local cluster for analysis. BISHOP operates similarly for the other sequential modes, *random* and *vector*, included in figure 1(a). By contrast, in *optimise* mode, a Bayesian optimisation model using Gaussian process (GP) regression [32, 57] is used to select simulation input parameters that are most likely to optimise the value of a target function. This results in the active feedback process of Bayesian optimisation, which is illustrated in figure 1(b) and presented in more detail in later sections.

In each simulation, the target front surface is irradiated at the focal position of the laser, defined as $X = 0 \mu\text{m}$. The time at which the peak of the pulse reaches $X = 0 \mu\text{m}$ is defined as $t = 0 \text{ fs}$. Laser energy is transferred to plasma electrons, generating a fast electron population that propagates through the target to drive the formation of the TNSA electrostatic sheath field at the rear side. This fast electron population has a broad exponential energy distribution of the form $\frac{dN}{dE} \propto \exp(-\frac{E_e}{k_B T_e})$, where E_e is the electron energy and T_e is the temperature of the population. The fast electron temperature increases with the peak laser irradiance [18–21] and is one of the key properties defining the maximum energy to which ions are accelerated via TNSA [15–17]. The fast electron temperature was determined for each simulation result by fitting to the high energy component of the fast electron distribution, sampled for energies $>0.511 \text{ MeV}$ up to 90% of the maximum electron energy. The lower bound ensures that the fit is made to relativistic electrons and the upper bound was selected to reduce the influence of statistical noise (which is high at the highest energies due to the low numbers of electrons). The coefficient of determination, R^2 , for the fit was ≥ 0.9 in all cases. Further details are available in the supplementary material file.

Figure 2(a) shows the fast electron temperature as a function of E_L and τ_L . For fixed τ_L , the laser intensity is increased by increasing E_L . A power scaling of the form $k_B T_e = a \cdot I_L^b$ with $b = 0.58 \pm 0.03$ is found, in good agreement with previously published results [18]. The maximum temperature is obtained for $E_L = 15 \text{ J}$ and $\tau_L = 33 \text{ fs}$. The same input parameters produce the maximum proton energy ($\epsilon_{pmax} = 109 \text{ MeV}$), as shown in figure 2(b). This is consistent with a strong overall correlation between ϵ_{pmax} and $k_B T_e$, and ϵ_{pmax} and E_L (thus I_L), as observed in figures 2(c) and (d).



The reduced dimensionality of 2D PIC simulations results in an overestimation of $k_B T_e$ [43] and ϵ_{pmax} [44] compared to 3D PIC simulations and experiment. A factor of 2.2 difference in ϵ_{pmax} is expected between 2D and 3D PIC simulations [44, 45]. To enable a comparison with experimental results, we include results from Dover *et al* [46] in figure 2(d) scaled up by this same factor. These experimental results were obtained for similar laser conditions as used in the present study; I_L in the range $(0.3 \rightarrow 3.4) \times 10^{21} \text{ W cm}^{-2}$, $\phi_L \approx 1.5 \mu\text{m}$, $\lambda_L \approx 800 \text{ nm}$ and $\tau_L \approx 40 \text{ fs}$.

Although the simulated ϵ_{pmax} values are higher, the observed power law scaling is similar to these example experimental results, and is generally in good agreement with a range of previously published TNSA scaling results [15–17, 36, 47]. As such, the simulation results provide a useful control data-set to evaluate the performance of a Bayesian optimisation algorithm in identifying optimal laser and target conditions for maximising proton energy.

2.2. Application of Bayesian optimisation

The results shown in figure 2 represent a conventional linear grid search of two key laser–target input variables that are known to influence the maximum proton energy achieved during an interaction. Optimal input conditions were identified for each of these variables, though many more parameters are known to influence source properties in the TNSA regime, such as the laser focal spot size [46, 48, 49], target thickness [50–52] and pre-plasma density scale length [53, 55]. As such, consideration of more interaction parameters will enhance the ability to optimise these novel ion sources. However, linear variation of more than two input parameters would require a substantial increase in the number of simulations required (an exponential growth with the number of parameters), which quickly becomes unfeasible given the computational resources required for high resolution simulations and the simulation times involved (each simulation requiring 1500 core-hours for the results presented in figure 2). To address this, an open source Bayesian optimisation algorithm [56] was integrated into the BISHOP code and used to identify the optimal conditions.

Bayesian optimisation [32, 57] is well suited to multidimensional problems and, as stated earlier, is particularly effective when the number of evaluations of the target function to be optimised are limited. Some other ML based optimisation techniques, such as genetic algorithms, require significantly less computational resource to generate a model of the target function, but require many more evaluations to identify an optimum. An example Bayesian optimisation procedure is illustrated in figure 1(b), for a simple

1D target function. The process begins by randomly generating a number of input values, x_n , which are evaluated by the target function, $f(x)$. There is equal weighting for the generation of any of these input values across the parameter space. A model is then produced to approximate the target function, by integrating over all possible functions which could fit the evaluations, $f(x_n)$, that have been made. In this study, GP regression is used to create the model [57]. A simplified version of this GP model, known as the acquisition function, is then constructed. This acquisition function is simple to evaluate compared to the true target function, meaning thousands of potential input values can be quickly evaluated. The candidate input value that provides the maximum value when evaluated by the acquisition function is chosen to be evaluated by the true target function. This process is demonstrated in figure 1(b), where the blue line represents the acquisition function, which is maximised at the position of the blue circle. This input value is then evaluated by the true target function as shown in figure 1(a) and the GP model is updated. The process is repeated, with the GP model approximating the true target function with increasing accuracy as more evaluations are made. This is shown in figure 1(b), where the variance of the GP model reduces as more data is collected and the target function is optimised within ten iterations.

In order to validate this approach, an open source Bayesian optimisation algorithm [56] was integrated into the BISHOP code and tasked with optimising the maximum proton energy over the same parameter space covered by the conventional linear grid search of figure 2. For this first optimisation exercise, the Bayesian optimisation algorithm was restricted to sampling parameter combinations that had been simulated as part of the data-set in figure 2 ($E_L = [3.0, 4.3, \dots, 15.0]$ J; $\tau_L = [25.0, 33.3, \dots, 100.0]$ fs) rather than having the freedom to select any input values between the bounds of $E_L = (3 \rightarrow 15)$ J and $\tau_L = (25 \rightarrow 100)$ fs.

The algorithm began by selecting input laser conditions at random and extracted ϵ_{pmax} for each simulation run. A GP model and acquisition function were then constructed and used to guide subsequent simulation inputs, until the optimal conditions from the linear grid search were identified. The initial random points had equal weighting across the parameter space and the boundaries were defined by the upper and lower values of the range over which the parameter space was evaluated. These boundaries are rectangular for two parameters and hyper-rectangular for higher dimensions. Given the shape of the parameter space for the TNSA mechanism, we would expect the extremes of our maximum proton energies to lie close to the boundaries. This poses no particular issue for the Bayesian optimisation approach.

In applying Bayesian optimisation to TNSA, we assume that the function describing the proton maximum energy over the parameter space is smooth and continuous. We also assume that the noise present in the parameter space does not obscure the function describing proton maximum energy and that this noise has constant variance. These conditions are well met when optimising TNSA with PIC simulations, but may prove harder to meet when applied to other acceleration mechanisms or to experimental data.

The optimisation was repeated for several different configurations of the Bayesian algorithm to find a setup that consistently identified the true optimum from the linear grid search in the smallest number of simulations. This involved varying the number of initial simulations with randomised inputs, as well as the manner in which the acquisition function guides input values throughout the optimisation process. Configurations were initialised with 5, 10 and 20 simulations of randomised inputs and the convergence of each to the true optimum from the linear grid search is shown in figure 3(a). After the randomised simulations are completed, the acquisition function guides the input conditions to be simulated and has a strong influence on the efficiency of the algorithm in identifying the optimum conditions.

Several acquisition functions are available, and are constructed in slightly different ways, with the upper confidence bound (UCB), probability of improvement (PI), and expected improvement (EI) being the most commonly used [32]. Typically, only one of these is used throughout the entire optimisation process, and the convergence of the algorithm to the true optimum of the target function may depend on which is selected. In this study, the algorithm samples candidate input values from each of these acquisition functions at every iteration, rather than just using one throughout, referred to as *hedging* [58]. This technique has been shown to maximise the performance of a Bayesian optimisation algorithm and negates the need to choose an optimal acquisition function from the outset [58]. Hedging increases the computational time required to select new input values between iterations of the optimisation process, but the increase is negligible when compared to the computational load required to perform each simulation.

When an acquisition function is selected, candidate input values are evaluated using the mean, μ_{GP} , and variance, $\sigma_{GP}(x)$, of the GP model. Each acquisition function contains a parameter that determines how much weight should be given to the mean and variance, respectively. For example, the UCB is constructed as $UCB(x) = \mu_{GP} + \kappa\sigma_{GP}(x)$, where a low κ value will cause the acquisition function to preferentially select input values in areas of the parameter space where the mean of the GP model is high, while penalising input values in areas of the parameter space where the model variance is high. In our case, this results in the

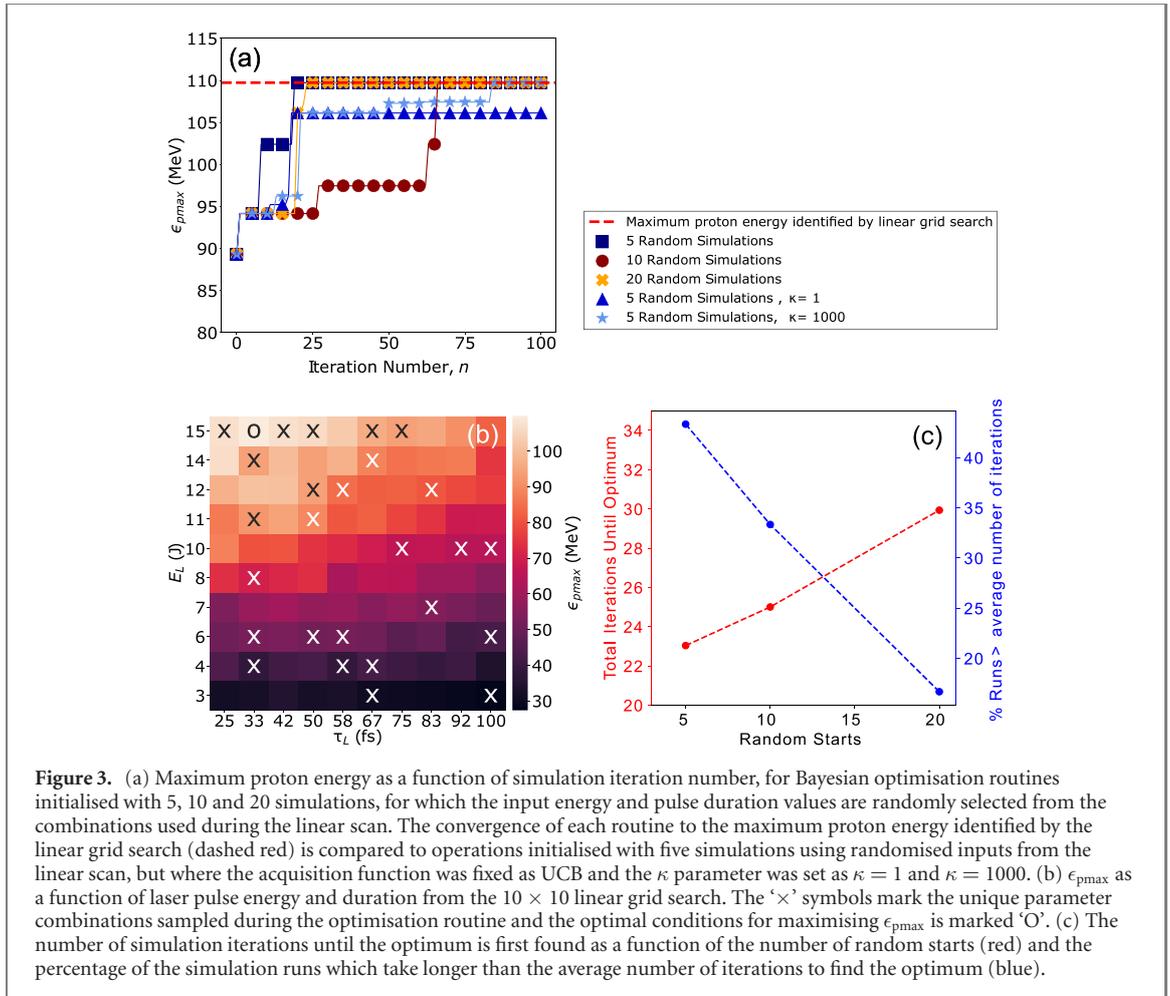


Figure 3. (a) Maximum proton energy as a function of simulation iteration number, for Bayesian optimisation routines initialised with 5, 10 and 20 simulations, for which the input energy and pulse duration values are randomly selected from the combinations used during the linear scan. The convergence of each routine to the maximum proton energy identified by the linear grid search (dashed red) is compared to operations initialised with five simulations using randomised inputs from the linear scan, but where the acquisition function was fixed as UCB and the κ parameter was set as $\kappa = 1$ and $\kappa = 1000$. (b) ϵ_{pmax} as a function of laser pulse energy and duration from the 10×10 linear grid search. The 'x' symbols mark the unique parameter combinations sampled during the optimisation routine and the optimal conditions for maximising ϵ_{pmax} is marked 'O'. (c) The number of simulation iterations until the optimum is first found as a function of the number of random starts (red) and the percentage of the simulation runs which take longer than the average number of iterations to find the optimum (blue).

algorithm running simulations with input values in areas of the parameter space that are well known to the model (uncertainty is low), and ϵ_{pmax} is predicted to be high, rather than exploring areas of the parameter space that are less well known to the model (uncertainty is high) because few simulations have been performed with input values in this region. The influence of κ is shown in figure 3(a), where the convergence of operations that solely used the UCB process are shown. The other acquisition functions mentioned, EI and PI, contain a parameter, ξ , which is analogous to κ in UCB and works in a similar fashion. In our case, where we apply a hedging strategy for optimising the acquisition function choice, values $\kappa = 1.96$ for UCB and $\xi = 0.01$ for PI/EI are used, similar to that used in reference [58].

The BISHOP code converged most quickly when initialised with 20 simulations of randomised inputs and when hedging was used to select the acquisition function, identifying the true optimum from the linear grid search ($\epsilon_{pmax} = 109$ MeV, at $E_L = 15$ J and $\tau_L = 33$ fs) in only 23 total simulations, as shown in figure 3(a). This example case corresponds to a reduction of ~ 28 h in the time required to identify these optimal conditions for the high performance computer used in this study, and a ~ 4 times decrease in the computational resources ($\sim 42\,000$ core-hours) required, compared to a systematic/linear grid search over 100 samples. Figure 3(b) shows the unique parameter combinations sampled during the optimisation routine on the full linear search grid.

The effect of the κ parameter was investigated over the range $\kappa = 1$ to $\kappa = 1000$. For $\kappa = 1$, the algorithm focuses too heavily on areas of the parameter space that are well known and where high ϵ_{pmax} is expected, but neglects to fully explore the parameter space, meaning it never identifies the true optimum. For $\kappa = 1000$, the algorithm spends too many iterations exploring the parameter space and does not focus heavily enough on simulating input values that are predicted to result in high ϵ_{pmax} , thus not identifying the true optima until very late in the process. This highlights the need to carefully consider the configuration of a Bayesian optimisation algorithm before applying it to a target function, though these considerations are reduced by hedging.

We further explore the sensitivity to the number of initial random runs by initialising the same optimisation for 5, 10 and 20 random simulations. Each optimisation run is repeated 30 times and allowed to run for 100 iterations. Figure 3(c) shows an average of the total number of iterations needed to first

identify the optimum value from the linear grid search as a function of the number of random starts. This shows that with more initial random simulations, the total number of simulations needed to find an optimum increases, but the number needed after the random simulations is reduced. Figure 3(c) also shows the percentage of the runs that take longer than the average number of iterations to find the optimum across the 30 run sample. As the number of initial random simulations is increased, there is more than a factor of two reduction in the number of simulation runs that take longer than average to identify the optimum value, leading to more consistent results. A balance must therefore be struck between increasing the number of starting simulations to improve consistency and limiting the total number of simulations that need to be run. The simulations detailed in the later sections of this article all use 20 initial random simulations for this reason.

2.3. Optimisation of proton acceleration with four input variables

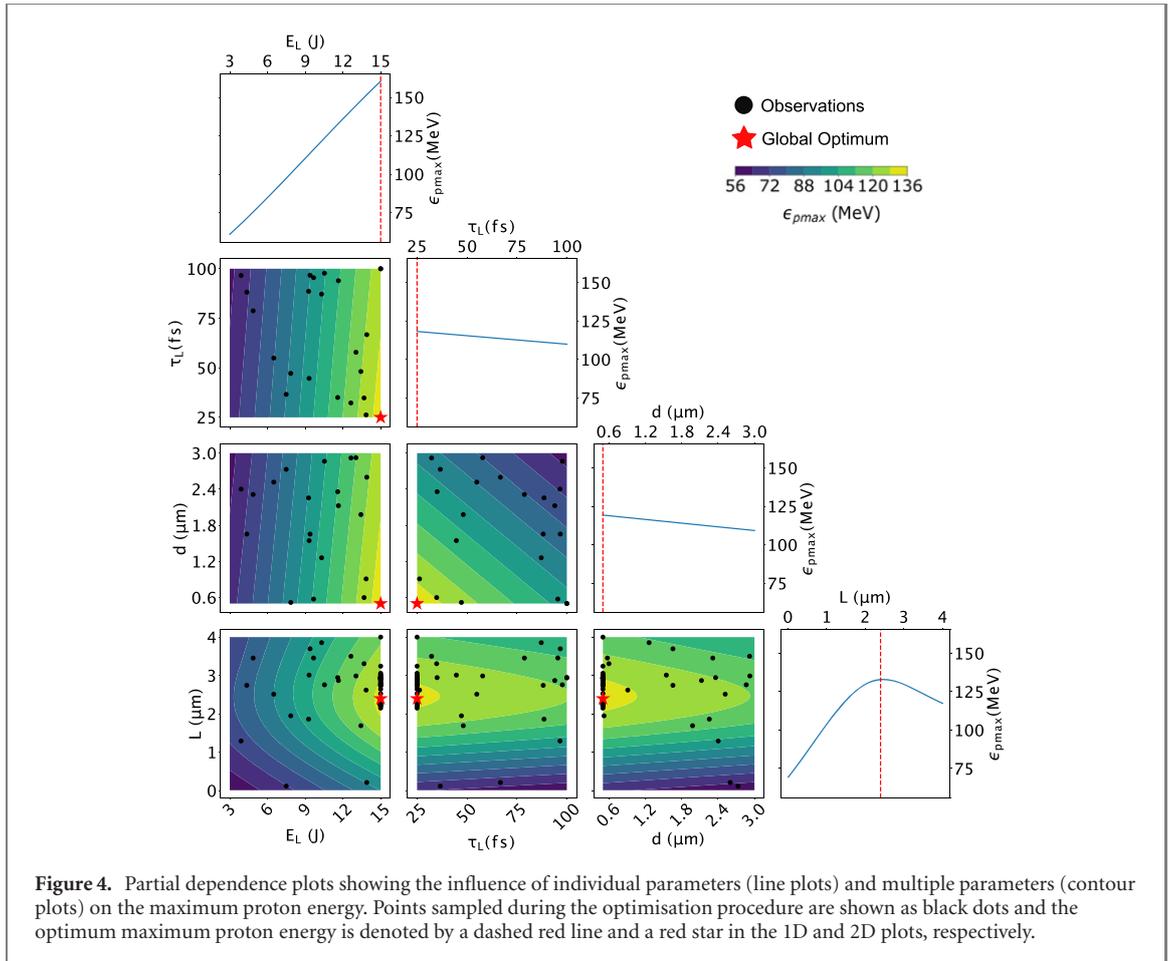
Having successfully applied Bayesian optimisation to achieve the maximum ϵ_{pmax} as a function of two input laser parameters, the technique was extended to include other important input variables. Among the most important for optimising TNSA are the target thickness, d [52], and the density scale length L at the target front side, which defines the pre-plasma density profile in the X -axis, described by $n_e(X) = n_{e0} \exp(X/L)$ for $X \leq 0$, where n_{e0} is solid density. The pre-plasma density profile influences the maximum proton energy by affecting laser energy absorption and the properties of the laser pulse via processes such as self-focusing or the onset of filamentation [53] (similar behaviour can also be produced via the addition of a low density layer to the front of a foil target [54]). The pre-plasma density profile can directly influence the coupling of laser energy into fast electrons [55] and induce additional heating mechanisms such as direct laser acceleration [59] or electron trapping [60]. Both d and L were therefore included as tunable parameters, together with E_L and τ_L , in a Bayesian optimisation procedure.

The BISHOP code was provided with the following boundary conditions: $E_L = (0.3 \rightarrow 15.0)$ J; $\tau_L = (25 \rightarrow 100)$ fs; $d = (500 \rightarrow 3000)$ nm; and, $L = (0.01 \rightarrow 4.00)$ μm . The input configurations sampled during this process are shown in the contour plots of figure 4 and the influence of individual input parameters on the maximum proton energy is shown in the line plots. After 20 simulations with randomised inputs and a further 30 simulations guided by a GP model, the algorithm converged to optimal conditions of $E_L = 15$ J, $\tau_L = 25$ fs, $d = 500$ nm and $L = 2.4$ μm , as shown in figure 4. This optimal laser–target configuration resulted in $\epsilon_{\text{pmax}} = 220$ MeV, which represents a twofold improvement over the results shown in figure 2. This highlights the importance of optimising for many input parameters, which is not always possible in a conventional grid search. In this case, extending the grid search of figure 2 to include d and L would require 10 000 simulations, running for over 10 000 h, which is prohibitively expensive.

Increasing E_L and reducing τ_L is shown to optimise ϵ_{pmax} in this TNSA-dominant regime, highlighting the expected dependency on peak laser intensity [15–17, 21]. This is expected for the laser parameter ranges explored.

Previous studies have shown that there is an optimum foil thickness for enhancing ϵ_{pmax} in the TNSA regime and that this is defined by a trade off between achieving efficient electron transport to the target rear side and minimising the influence of target expansion, arising from preheating driven by the rising edge of the laser pulse, which reduces the magnitude of the electric field responsible for ion acceleration [16, 51, 52]. ϵ_{pmax} is maximised for the thinnest target in this example study (500 nm), which is expected for the parameter range explored. For this demonstration, we have avoided ultrathin foils, for which other ion acceleration mechanisms (RPA and transparency-enhanced processes) compete with TNSA, so as not to complicate this first demonstration of the approach. The reasons for the observed optimum pre-plasma density scale length observed in figure 4 ($L = 2.4$ μm) are explored in the next section.

For the optimum that is identified here (to maximise ϵ_{pmax}), we are confident that it represents the true global optimum within our parameter space. After the first 20 simulations to generate the model, E_L and d were not varied by the model, and τ_L was only varied twice in the subsequent 60 iterations before the algorithm was manually stopped. The model quickly identified the expected optimisation for these parameters, which is the thinnest target, the lowest τ_L and the highest E_L . This is consistent with the grid scan results of figure 2 and our understanding of the TNSA mechanism (the ion energy scales with laser pulse intensity). Between iteration 20 and 50, L was tuned between 2.5 μm and 4 μm , and after the optimum of $L = 2.4$ μm was identified, the model only varied L within 5% of this value, suggesting a high degree of confidence. Identifying true global optima in a given multidimensional parameter space is a challenge, especially when evaluation of the parameter space is expensive, because not all points can be sampled. Any optimisation process is at risk of identifying local optima within the parameter space, while failing to identify the true global optima. In our study, this risk has been mitigated by testing multiple configurations of the algorithm on the grid search data of figure 2. A configuration that consistently

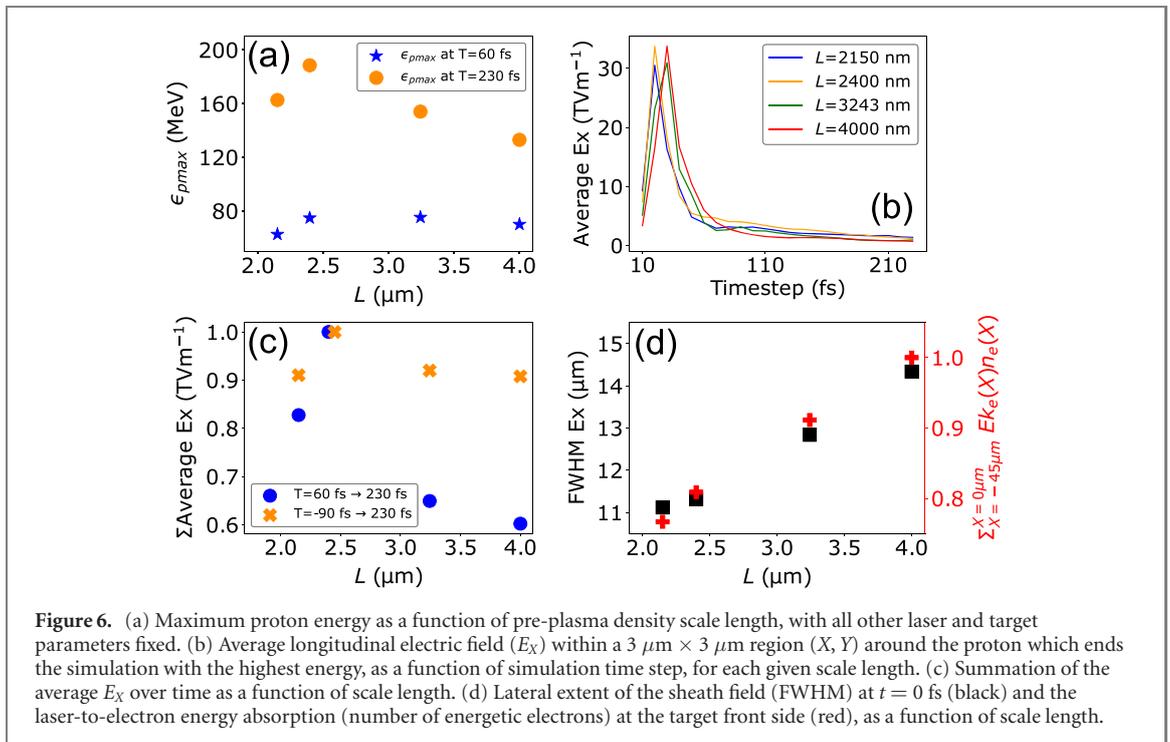
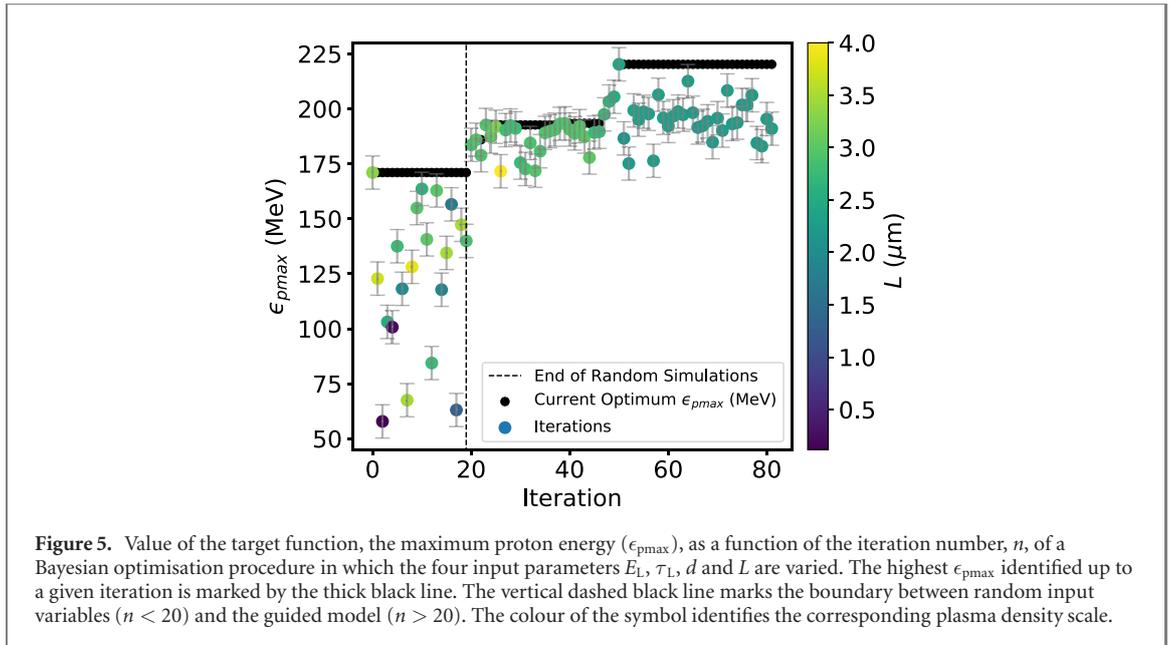


identified the true global optimum from the grid search was used for the four parameter optimisation case, for which a grid search is not possible (computationally prohibitive) and the true global optimum is not known in advance.

2.4. Exploration of the optimum parameters

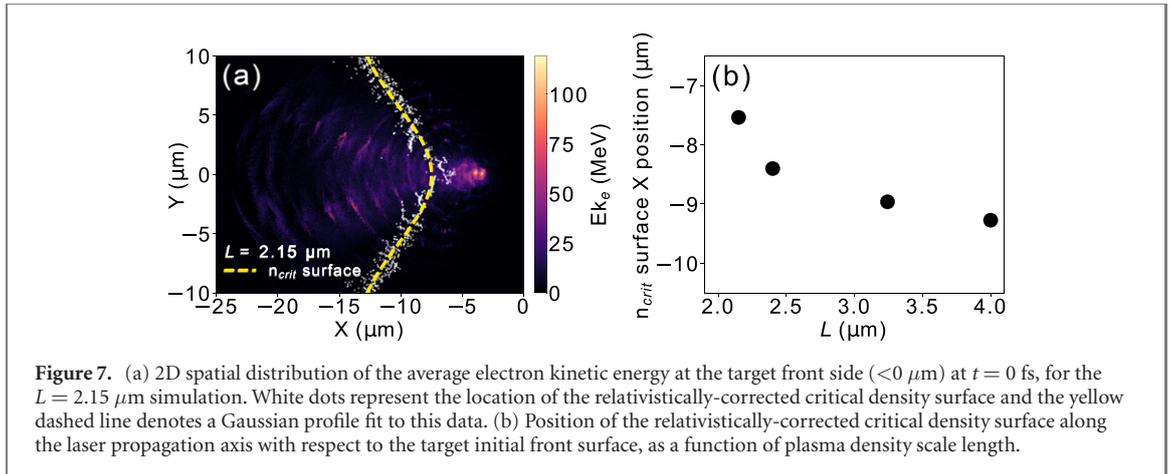
To explore the optimisation results in more detail, we review the input parameter values sampled at each iteration after the initial 20 simulations with randomised inputs were completed. Thereafter, the input parameter values were guided by the GP model, and an immediate improvement in ϵ_{pmax} is observed after the 20th iteration, as shown in figure 5. The values of E_L , τ_L and d selected by the algorithm at this iteration were those which proved to be optimal by the end of the optimisation routine ($E_L = 15$ J, $\tau_L = 25$ fs and $d = 500$ nm). On subsequent iterations, the algorithm mostly varied the pre-plasma scale length, while keeping the values of the other parameters constant. Initially, the algorithm explored relatively long scale lengths in the range 2.8–4.0 μm , before finding an optimum at iteration number 50 ($L = 2.4$ μm , resulting in $\epsilon_{pmax} = 220$ MeV). From this iteration until the operation was manually terminated, the algorithm explored shorter scale lengths in the range 2.1–2.4 μm . Investigating these results in more detail, small scale length fluctuations were noted to induce significant changes in ϵ_{pmax} . As such, a repeat simulation was performed at the optimal laser and target parameters with the random seed enabled in EPOCH [39]. The initial location and momentum distribution of the macroparticles in EPOCH is generated using a random number generator. A seed number is used to initialise the generator which can be fixed or random. With the random seed on, the maximum proton energy was reduced by 13%, from 220 MeV to 191 MeV on this repeat simulation, suggesting that statistical fluctuations contribute to changes in the precise proton energy produced for any given L .

To ensure that the scale length optimisation observed in these results is not driven by noise, ϵ_{pmax} was sampled for 13 simulations for which L was within 5% of the optimum value (average $L = (2.31 \pm 0.03)$ μm). The average ϵ_{pmax} within this batch of simulations was 200 MeV and the standard deviation was 8 MeV, which defines the error-bars on the data in figure 5. This is compared to 20 simulations with average $L = (2.91 \pm 0.06)$ μm , for which $\epsilon_{pmax} = (187 \pm 6)$ MeV. Thus, the variation of ϵ_{pmax} with L is statistically significant when compared to the variation produced by statistical fluctuations.



Previous studies have shown that, compared to the case of a sharp density step, ϵ_{pmax} can be enhanced if significant self-focusing of the laser pulse occurs when propagating within the pre-plasma and that there is an optimum density scale length for the enhancement [53]. We see evidence of self-focusing occurring in the simulation results presented in figure 6, which increases the laser intensity over the nominal value. However the degree of self-focusing does not change significantly over the scale length range explored and therefore this phenomena does not explain the optimisation of ϵ_{pmax} —see the supplementary material for further details.

To understand why ϵ_{pmax} is optimised at $L = 2.4 \mu\text{m}$, four input configurations sampled during the optimisation routine were repeated with particle tracking enabled in EPOCH such that the source and trajectory of the highest energy protons could be investigated. These repeat simulations were run for a reduced simulation time, finishing at time $t = 230$ fs, compared to $t = 560$ fs in the optimisation routine, due to the increased computational resource required to include particle outputs. The optimal input conditions of $E_L = 15$ J, $\tau_L = 25$ fs and $d = 500$ nm were fixed and L was varied from $2.15 \mu\text{m}$ to $4.00 \mu\text{m}$.



The maximum proton energy at $t = 230$ fs is found to vary with L as shown in figure 6(a), with an optimum at $L = 2.40 \mu\text{m}$, in good agreement with the results of the Bayesian optimisation routine. The highest energy proton at the end of the simulation was identified and tracked over all time in each of the four cases considered. The longitudinal electric field strength (E_X) was averaged within a $3 \mu\text{m} \times 3 \mu\text{m}$ region (X, Y) around the tracked proton at each time. In figure 6(b), the tracked proton is shown to experience a peak electric field at $t \sim 25$ fs, corresponding to the time at which hot electrons reach the target rear surface. This peak field occurs slightly later for the larger L cases, for which the electrons are generated further from the initial target front surface and propagate through an extended longitudinal plasma before reaching the target rear.

After the tracked proton leaves the target rear surface (at $t \sim 60$ fs), it experiences a higher electric field over the remaining time for the two shortest scale length cases, as shown in figure 6(c), where the average electric field is summed from $t = -90$ fs to $t = 230$ fs, and from $t = 60$ fs to $t = 230$ fs. The tracked proton experiences a stronger total field from $t = -90$ fs to $t = 230$ fs in the longer scale length simulations compared to $L = 2.15 \mu\text{m}$, due to the temporally broader peak field experienced by the proton (see figure 6(b)). This broader peak field results from an increased laser-to-electron energy absorption in the case of the longer scale length plasma, as shown in figure 6(d), where absorption is approximated by the electron energy obtained by multiplying the fast electron density and average kinetic energy in the region at the front side of the target ($X = [-45 \rightarrow 0] \mu\text{m}$). These factors contribute to a higher ϵ_{pmax} at $t = 60$ fs for longer scale lengths compared to $L = 2.15 \mu\text{m}$, as shown in figure 6(a).

Although a longer scale length plasma density profile at the target front side results in increased laser-to-electron energy absorption and a stronger electric field at early time steps, it also produces a spatially broader sheath field at the target rear surface. This is quantified by summing the E_X field for each cell along the Y -axis within $3 \mu\text{m} \times 40 \mu\text{m}$ (X, Y) of the target rear and fitting a Gaussian profile. The full width at half maximum (FWHM) of this fitted profile is shown to increase with L in figure 6(d). The larger sheath area results from the source of the electrons at the target front side shifting further from the target rear surface due to the expanded pre-plasma.

To quantify this, the position of the relativistically-corrected critical density surface, γn_{crit} , where γ is the electron Lorentz factor, is determined at each point along the Y -axis by averaging the electron kinetic energy over all cells from $X = [-45 \rightarrow 0] \mu\text{m}$ (i.e. on the front side of the target) and at $t = 0$ fs. The position of γn_{crit} is defined as the first X position where $n_e/\gamma n_{\text{crit}} = 1$, and is shown at each point in the Y plane for $L = 2.15 \mu\text{m}$ in figure 7(a). The γn_{crit} position along the laser propagation axis is determined from the peak of a Gaussian fit to this data and, as shown in figure 7(b), is found to be located farther from the target front surface as L is increased. The tracked proton propagates within $\pm 2 \mu\text{m}$ of the target normal axis in all simulations and thus experiences a weaker accelerating field for longer scale lengths, despite higher laser-to-electron energy absorption at the front surface. This explains the results of figure 6(a), where the maximum proton energy is initially ($t = 60$ fs) lower for $L = 2.15 \mu\text{m}$ compared to the larger L cases, but is higher by the end of the simulation at $t = 230$ fs. As such, the optimisation of ϵ_{pmax} with L is found to be the result of competition between achieving high laser-to-electron energy absorption at the target front side, which defines the fast electron number and temperature, and the sheath area at the rear side, which defines the fast electron density. Both the electron temperature and density are important parameters in defining the maximum energy to which ions are accelerated in the sheath expansion [15].

3. Summary

In summary, we report on the optimisation of laser-driven proton acceleration in PIC simulations of intense laser–plasma interactions using a Bayesian algorithm to efficiently identify the optimum laser and target parameters. A new code, called BISHOP, has been developed as a platform to automatically generate and process PIC simulation data with minimal user input, simplifying and reducing the time taken to undertake large scale simulation studies involving systematic scanning of parameters. The energy and duration of a high power laser pulse incident on a 1 μm -thick planar CH target was varied over 100 2D PIC simulations and optimum parameters for maximising proton energy in the TNSA-dominant regime were identified, to demonstrate this automated-PIC parameter-scanning approach.

We next incorporated a Bayesian optimisation algorithm into the BISHOP code to create an efficient method of optimising a selected output parameter as a function of multiple input parameters. The viability and efficiency of this approach was demonstrated via an investigation of the influence of several laser and target input parameters on the maximum energy to which protons are accelerated by the TNSA mechanism. Using this approach, the optimum conditions were identified four times faster than by varying two laser parameters (pulse energy and duration) linearly over 100 2D simulations. In the case of four input parameters (including target thickness and front-side density scale length) the optimum conditions are identified 1000 times quicker compared to systematic linear scanning. This four-parameter optimisation study resulted in a doubling of the maximum proton energy compared to the two-parameter test, highlighting the potential to use this approach to identify combinations of parameters that enhance a desired beam property.

Our investigation also identifies an optimum pre-plasma density scale length for achieving high proton energies, which results from balancing total laser energy coupling to fast electrons with the electron density in the sheath. This optimal scale length configuration is not obvious from previously published results and may not have been identified without a ML technique due to the prohibitively high cost of performing a conventional linear grid search over four input parameters at the resolution needed. Extension of this approach to a larger number of input parameters is likely to identify new correlations with output beam properties, inter-dependencies and conditions for simultaneously optimising several desired output beam properties.

This study is a demonstration of Bayesian optimisation applied to laser-driven ion acceleration, using 2D PIC simulations. This can be extended in future work to include more computationally demanding 3D PIC simulations for quantitative predictions. The Bayesian optimisation approach adopted here can also be extended to the generation of other types of particles and radiation in numerical investigations of laser–plasma interactions. The BISHOP code can be developed to apply various other ML algorithms, both to directly optimise simulation conditions and to benchmark the suitability of these algorithms for application in experimental investigation of laser–plasma interactions.

Acknowledgments

This work is financially supported by EPSRC (Grant Numbers EP/R006202/1 and EP/V049232/1) and STFC (Grant Number ST/V001612/1). It used the ARCHIE-WeSt and ARCHER2 high performance computers, with access to the latter provided via the Plasma Physics HEC Consortia (EP/R029148/1), and the University of Cambridge Research Computing Service (funded by Grant Number EP/P020259/1). EPOCH was developed under EPSRC Grant EP/G054940/1. The research has also received funding from Laserlab-Europe (Grant Agreement No. 871124, European Union's Horizon 2020 research and innovation programme).

Data availability statement

The data that support the findings of this study will be openly available following an embargo at the following URL/DOI: <https://doi.org/10.15129/3ffd73de-cf0c-4a14-906f-05cd6a79077e>.

ORCID iDs

E J Dolier  <https://orcid.org/0000-0002-2744-7406>

R J Gray  <https://orcid.org/0000-0003-0610-9595>

P McKenna  <https://orcid.org/0000-0001-8061-7091>

References

- [1] Daido H, Nishiuchi M and Pirozhkov A S 2012 *Rep. Prog. Phys.* **75** 056401
- [2] Macchi A, Borghesi M and Passoni M 2013 *Rev. Mod. Phys.* **85** 751–93
- [3] Bulanov S V, Esirkepov T Z, Khoroshkov V S, Kuznetsov A V and Pegoraro F 2002 *Phys. Lett. A* **299** 240
- [4] Bulanov S V, Wilkens J J, Esirkepov T Z, Korn G, Kraft G, Kraft S D, Molls M and Khoroshkov V S 2014 *Phys.-Usp.* **57** 1149
- [5] Zeil K et al 2013 *Appl. Phys. B* **110** 437–44
- [6] Sun Z 2021 *AIP Adv.* **11** 040701
- [7] Kroll F et al 2022 *Nat. Phys.* **18** 316–22
- [8] Dromey B et al 2016 *Nat. Commun.* **7** 10642–7
- [9] Barberio M et al 2017 *Sci. Rep.* **7** 12009
- [10] Barberio M et al 2018 *Nat. Commun.* **9** 372
- [11] Roth M et al 2001 *Phys. Rev. Lett.* **86** 436–9
- [12] Fernández J C, Albright B J, Beg F N, Foord M E, Hegelich B M, Honrubia J J, Roth M, Stephens R B and Yin L 2014 *Nucl. Fusion* **54** 054006
- [13] Snavely R A et al 2000 *Phys. Rev. Lett.* **85** 2945
- [14] Wilks S C et al 2001 *Phys. Plasmas* **69** 1383
- [15] Mora P 2003 *Phys. Rev. Lett.* **90** 185002
- [16] Fuchs J et al 2006 *Nat. Phys.* **2** 48–54
- [17] Robson L et al 2007 *Nat. Phys.* **25** 58–62
- [18] Wilks S C, Kruer W L, Tabak M and Langdon A B 1992 *Phys. Rev. Lett.* **69** 1383
- [19] Beg F N et al 1997 *Phys. Plasmas* **4** 447
- [20] Haines M G et al 2009 *Phys. Rev. Lett.* **102** 045008
- [21] Kluge T et al 2011 *Phys. Rev. Lett.* **107** 205003
- [22] Esirkepov T, Borghesi M, Bulanov S V, Mourou G and Tajima T 2004 *Phys. Rev. Lett.* **92** 175003
- [23] Robinson A P L, Gibbon P, Zepf M, Kar S, Evans R G and Bellei C 2009 *Plasma Phys. Control. Fusion* **51** 024004
- [24] Palmer C A J et al 2011 *Phys. Rev. Lett.* **106** 014801
- [25] Kar S et al 2012 *Phys. Rev. Lett.* **109** 185006
- [26] Macchi A, Veghini S and Pegoraro F 2009 *Phys. Rev. Lett.* **103** 085003
- [27] McIlvenny A et al 2020 *Plasma Phys. Control. Fusion* **62** 054001
- [28] Yin L, Albright B J, Hegelich B M, Bowers K J, Flippo K A, Kwan T J T and Fernández J C 2007 *Phys. Plasmas* **14** 056706
- [29] Higginson A et al 2018 *Nat. Commun.* **9** 724
- [30] Brantov A V et al 2016 *Phys. Rev. Lett.* **116** 085004
- [31] Noaman-ul-Haq M et al 2017 *Phys. Rev. Accel. Beams* **20** 041301
- [32] Shahriari B et al 2015 *Proc. IEEE* **103** 148
- [33] Shaloo R J et al 2020 *Nat. Commun.* **11** 6355
- [34] Dann S J et al 2019 *Phys. Rev. Lett.* **22** 041303
- [35] Jalias S, Kirchen M, Messner P, Winkler P, Hübner L, Dirkwinkel J, Schnepf M, Lehe R and Maier A R 2021 *Phys. Rev. Lett.* **126** 104801
- [36] Takagi Y et al 2021 *Phys. Rev. Res.* **3** 043140
- [37] Djordjević B Z et al 2021 *Phys. Plasmas* **28** 043105
- [38] Smith J R, Orban C, Morrison J T, George K M, Ngirmang G K, Chowdhury E A and Roquemore W M 2020 *New J. Phys.* **22** 103067
- [39] Arber T D et al 2015 *Plasma Phys. Control. Fusion* **57** 113001
- [40] Danson C, Hillier D, Hopps N and Neely D 2015 *High Power Laser Sci. Eng.* **3** e3
- [41] Kühn S et al 2017 *J. Phys. B: At. Mol. Opt. Phys.* **50** 132002
- [42] Zou J P et al 2015 *High Power Laser Sci. Eng.* **3** e2
- [43] Stark D J, Yin L, Albright B J and Guo F 2017 *Phys. Plasmas* **24** 053103
- [44] d’Humières E et al 2013 *Phys. Plasmas* **20** 023103
- [45] Sgattoni A et al 2020 *Phys. Rev. E* **85** 036405
- [46] Dover N P et al 2020 *Phys. Rev. Lett.* **124** 084802
- [47] Zimmer M et al 2021 *Phys. Rev. E* **104** 045210
- [48] Coury M et al 2012 *Appl. Phys. Lett.* **100** 074105
- [49] Brenner C M, McKenna P and Neely D 2014 *Plasma Phys. Control. Fusion* **56** 084003
- [50] Mackinnon A J et al 2002 *Phys. Rev. Lett.* **88** 215006
- [51] Kaluza M et al 2004 *Phys. Rev. Lett.* **93** 045003
- [52] Neely D, Foster P, Robinson A, Lindau F, Lundh O, Persson A, Wahlström C-G and McKenna P 2006 *Appl. Phys. Lett.* **89** 021502
- [53] McKenna P et al 2008 *Laser Part. Beams* **26** 591
- [54] Bin J H et al 2015 *Phys. Rev. Lett.* **115** 064801
- [55] Gray R J et al 2014 *New J. Phys.* **16** 113075
- [56] Pedrogosa F et al 2011 *J. Mach. Learn. Res.* **12** 2825–30
- [57] Rasmussen C E et al 2006 *Gaussian Processes for Machine Learning* (Cambridge, MA: MIT Press)
- [58] Hoffman M W et al 2011 Portfolio allocation for Bayesian optimization *Proc. Conf. Uncertainty Artif. Intell.* p 327
- [59] Arefiev A V, Khudik V N, Robinson A P L, Shvets G, Willingale L and Schollmeier M 2016 *Phys. Plasmas* **23** 056704
- [60] Lefebvre E and Bonnaud G 1997 *Phys. Rev. E* **55** 1011