



PAPER

OPEN ACCESS

RECEIVED

3 February 2026

ACCEPTED FOR PUBLICATION

26 March 2026

PUBLISHED

16 April 2026

Original content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

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



Bayesian optimisation of non-linear Breit–Wheeler pair production in simulated laser experiments

Christopher Arran^{1,2,*} , Stuart Morris³  and Christopher P Ridgers⁴ ¹ Physics Department, Lancaster University, Bailrigg LA1 4YW Lancaster, United Kingdom² Cockcroft Institute, Keckwick Ln, Daresbury, Warrington WA4 4AD, United Kingdom³ Department of Physics, University of Warwick, CV4 7AL Coventry, United Kingdom⁴ York Plasma Institute, University of York, Heslington, York YO10 5DD, United Kingdom

* Author to whom any correspondence should be addressed.

E-mail: c.arran@lancaster.ac.uk**Keywords:** strong field QED, laser wakefield acceleration, Monte Carlo methods, Breit–Wheeler process, inverse Compton scattering

Abstract

High laser intensities enable the production of electron–positron pairs from bright gamma rays passing through strong fields. Potentially the most promising approach for all-optical experiments in the near term uses dense but higher divergence electron beams from laser wakefield acceleration to produce gamma rays through inverse Compton scattering. Achieving many-photon collisions between these gamma rays and the high intensity laser pulse in practice is extremely difficult, however, due to significant shot-to-shot jitter in laser pointing and timing. In practice, this jitter reduces the yield of electron–positron pairs by orders of magnitude. We model these practical difficulties using simulated Monte–Carlo experiments. By using a more efficient algorithm for sampling infrequent pair production with particle splitting, we enable the exploration of a multi-dimensional parameter space. Using Gaussian Process Regression we then efficiently find optimal conditions for maximising pair production by changing the laser spot size, the energy in the colliding beam, and the stand-off distance between the laser wakefield accelerator and the focus of the colliding laser pulse. We find that the optimal stand-off distance increases with the degree of laser jitter and that the best conditions for producing electron–positron pairs are not the same as the best conditions for maximising the energy in the gamma rays. With 100 J of laser energy, we estimate realistic rates of pair production of around 1 pair per 100 electrons are achievable even with jitter of 10 s of microns and 10 s of femtoseconds.

1. Introduction

Laser facilities continue to push the intensity frontier to new records, in recent years achieving peak powers of 10 PW [1] and peak intensities of $10^{23} \text{ W cm}^{-2}$ [2]. Several studies have explored the possibility of creating matter from pure light in sufficiently strong laser fields [3–5]. In the centre-of-mass frame of this many-photon interaction, the field strength approaches the Schwinger field [6], enabling the production of positron–electron pairs in vacuum through the non-linear Breit–Wheeler process [7]. Previous experiments have been able to reach this regime using a 47 GeV electron beam from SLAC colliding with an 0.4 TW laser pulse to produce around 100 positron–electron pairs from around 10^{10} electrons [8]. This first produces bright γ -rays from inverse Compton scattering [9], before the γ -rays themselves interact with the laser field. Another approach has been to use the strong atomic fields present within periodic crystal lattices [10, 11], with a 200 GeV electron beam driving trident pair production in a germanium crystal. Now, several laser facilities aim to utilise unprecedented laser powers and produce positron–electron pairs in vacuum by using laser-driven plasma wakefield accelerators [12, 13].

Conducting these experiments in practice is significantly more complicated than may be imagined from idealised simulations, with the interactions requiring femtosecond- and micron-level synchronisation and alignment between the laser and electron beams. In order to successfully produce positron–electron pairs in vacuum in laser-driven experiments, this spatial and temporal jitter must be accounted for, typically by running many simulations over the range of possible configurations. However, running this many simulations is often computationally expensive, especially when the chances of producing a positron–electron pair can be vanishingly small with current-generation laser facilities. In order to determine the best experimental parameters from a large domain of options, it is necessary to increase the efficiency both of calculating the pair production rates, and of exploring a multi-dimensional parameter space.

Several Monte-Carlo and particle-in-cell (PIC) codes have already increased computational efficiency by introducing particle merging [14–16] and splitting [17–20] algorithms. Merging techniques are useful when lots of particles are being produced, using too much computational memory, whereas splitting techniques are useful when few particles are produced, allowing codes to explore sparse areas of the parameter space and reduce the variance in the results. With Breit–Wheeler pair production a rare process with current laser parameters, particle splitting is a promising approach. However, most particle splitting algorithms such as references [18, 20] split the primary particle (in our case photons), which we generally have many of already, unnecessarily increasing computation in the particle push or tracking. On the other hand, algorithms like references [17, 19] split the secondary particle (in our case a positron–electron pair) only after it has been produced, which is unhelpful if the chances of it being produced in the first place are so small. We therefore want to perform particle splitting during the pair production step, which is a novel and computationally efficient approach but risks affecting the algorithm’s accuracy.

Here, we find optimum parameters for electron–positron pair production in all-optical experiments by employing two techniques for performing many simulations with higher computational efficiency. First, we describe a particle splitting algorithm that performs the split during pair production, and show how it requires changes to the way we model pair production. This accurately calculates the rate of pair production for rare events without significantly increasing the runtime, allowing us to rapidly conduct many simulated experiments, including the effect of jitter. Second, we use a Bayesian optimisation scheme built on Gaussian process regression to efficiently explore the parameter space and find the maximum rate of pair production. This scheme is significantly faster than a brute force grid search and builds a fast surrogate model for estimating the rate of pair production across the domain. We show that more positron–electron pairs can be produced by combining the tightest possible laser focus with a few-cm-scale stand-off distance between the laser wakefield accelerator and the interaction point. This ameliorates the effect of jitter while retaining the high field strength required for pair production.

2. Pair production and particle splitting

2.1. Pair production algorithm

In both Monte-Carlo and PIC codes, the emission of high energy photons and the production of positron–electron pairs are treated as discrete random and independent events which can occur in any fixed time interval. As such, they can be modelled using a Poisson distribution, where the average number of emissions per particle in a given timestep δt are given by:

$$\lambda_\gamma = \frac{dN_\gamma(\chi_e)}{dt} \delta t \quad (1)$$

$$\lambda_{e^+,e^-} = \frac{dN_{e^+,e^-}(\chi_\gamma)}{dt} \delta t, \quad (2)$$

where the electron and photon quantum parameters are given by $\chi_e = |F_{\mu\nu} p_e^\nu|/m_e E_s$ and $\chi_\gamma = |F_{\mu\nu} \hbar k_\gamma^\nu|/m_e E_s$, for a Schwinger field $E_s = m_e^2 c^3 / e \hbar$. The probability of emitting n particles in a given timestep is then

$$P(n|\lambda) = \frac{\lambda^n e^{-\lambda}}{n!}. \quad (3)$$

Often, in order to reduce the computational expense of sampling from a Poisson distribution at every timestep, codes use an optical depth approach (for instance in the EPOCH PIC code [21, 22]).

If the emission rate is sufficiently low, the chances of two or more emissions in a single timestep is negligible and the probability of a single particle being emitted on a timestep before t_N becomes:

$$P_{\text{emission}}(t_0 < t < t_N) \approx 1 - P_{\text{no emission}}(t_0 < t < t_N) \\ \approx 1 - \exp\left[-\sum_{n=0}^N \lambda(t_n)\right].$$

Now a uniformly distributed random number in the range $0 \leq r < 1$ is chosen to represent $P(t)$ and an emission occurs on the first timestep when $1 - \exp\left[-\sum_n \lambda(t_n)\right] \leq r$, which is when the accumulated optical depth reaches $\sum_n \lambda(t_n) \geq -\ln(1-r)$. Only one random number must be chosen for each emission event, calculated the moment a particle is created or the moment after an emission takes place, and a single running total tracks the optical depth. The optical depth method therefore efficiently models rare emissions which occur singly and consecutively.

This model is particularly relevant for pair production, as a photon ceases to exist after emitting a positron–electron pair, so more than one emission is impossible. In fact, emitting a photon also changes an electron’s trajectory, due to radiation reaction, and so emitting multiple photons in a single timestep using the same value for χ_e is also unphysical if the photon energy is sufficiently high. The optical depth method is therefore generally very well suited for the problems we are modelling.

However, if emissions are very rare, such as for Breit–Wheeler pair production, a tremendously large number of particles must be used to model the likelihood of even a single emission, with the vast majority of these particles not participating in the interaction at all. Each of the particles must be tracked through the simulation even if they do not contribute at all to the creation of pairs. In this situation, despite the strengths of the optical depth method, modelling single rare emissions is extremely computationally inefficient.

2.2. Particle splitting algorithms

We therefore want to introduce particle splitting into the calculation of Breit–Wheeler pair production in Monte-Carlo and PIC codes, during the pair production step rather than before or afterwards. In PIC codes, each macroparticle represents a large number of underlying real particles, which move around the simulation grid together and sample the underlying distribution in (\mathbf{x}, \mathbf{p}) space with a given weight w . When PIC codes model random emission processes they typically use the method discussed in section 2.1, sampling the probability distribution function once per timestep and creating child particles which also each have a weight of w . Here we describe how we can let each macroparticle move around as before, but instead independently sample the chances of producing a positron–electron pair many times per timestep at a splitting number uw , making the average number of emissions per timestep now equal to:

$$\lambda_{e^+,e^-} \rightarrow uw \frac{dN_{e^+,e^-}(\chi_\gamma)}{dt} \delta t. \quad (4)$$

The weight of any pairs produced will be reduced proportionately, from w to $1/u$, and the weight of the photon will be slightly reduced to account for the decay, from w to $w - 1/u$. If this weight falls to zero the photon is removed from the simulation. Particle splitting does not affect the number of parent particles and does not significantly increase the cost of the particle push, only affecting the number of pair particles. In this way, emission events can be made more common without greatly increasing the runtime, more efficiently modelling pair production.

However, there is a risk to naively using the optical depth algorithm with a high rate of particle splitting for this purpose. Firstly, if the degree of particle splitting u is sufficiently high that $\lambda_{e^+,e^-} \sim 1$, there is now a non-negligible chance that more than one pair will be produced per timestep, giving $P_{\text{emission}}(0 < t < t_N) \neq 1 - \exp\left[-\sum_{n=1}^N \lambda(t_n)\right]$. Secondly, even when only one or fewer pairs is produced in a timestep, the optical depth is refreshed after each emission using a new random number. This wipes the history for all particles represented by the macroparticle, not just the fraction $1/u$ that has been emitted, making emissions no longer independent. Without particle splitting, neither of these are issues, as photons are killed immediately upon producing just one electron–positron pair. Using particle splitting while using a single optical depth per particle, however, can give a drastically different estimate for the rate of pair production.

This problem can be resolved in three different ways, with different degrees of accuracy:

I—Sampling the full Poisson distribution on every timestep to model particle splitting. The average number of emissions in each timestep is still increased by a factor u but now any number of emissions

Table 1. Summary of the different particle splitting techniques for an average number of pairs produced per timestep $\lambda = uw \frac{dN_{\pm}}{dt} \delta t$, including a particle splitting rate u . Columns describe the approximate computational cost (as the average number of random numbers required per macroparticle per timestep), the range of λ where the algorithm is appropriate, whether the algorithm keeps weights simple (either w or $1/u$), and the energy conservation of the algorithm.

<i>Algorithm</i>	<i>Computation</i>	<i>Valid when:</i>	<i>Simple weights?</i>	<i>Conserves energy?</i>
Naive optical depth	$\text{Min}(\lambda, 1)$	$\lambda \ll 1$	Yes	Exactly
Additive optical depth	$\text{Min}(\lambda, 1)$	$\lambda \lesssim 1$	Yes	On average
Additive plus subcycling	λ	Any λ	Yes	On average
Poisson	$\text{Max}(\lambda, 1)$	Any λ	No	Exactly

are possible in any timestep. If at least one emission occurs in a single timestep, $n > 0$, a single pair of electron and positron macroparticles can be produced, each with a weight n/u . Alternatively, if a larger number of pair particles are desired, for instance to better sample an energy distribution, n macroparticles can be produced, each with a weight $1/u$. In either case, the weight of the photon is reduced to $w - n/u$. This is a relatively simple extension for codes like PTARMIGAN [23], which already calculate the emission probability on every timestep. It ensures that pair production events are independent and conserves energy on every emission, while allowing for an arbitrarily large degree of particle splitting.

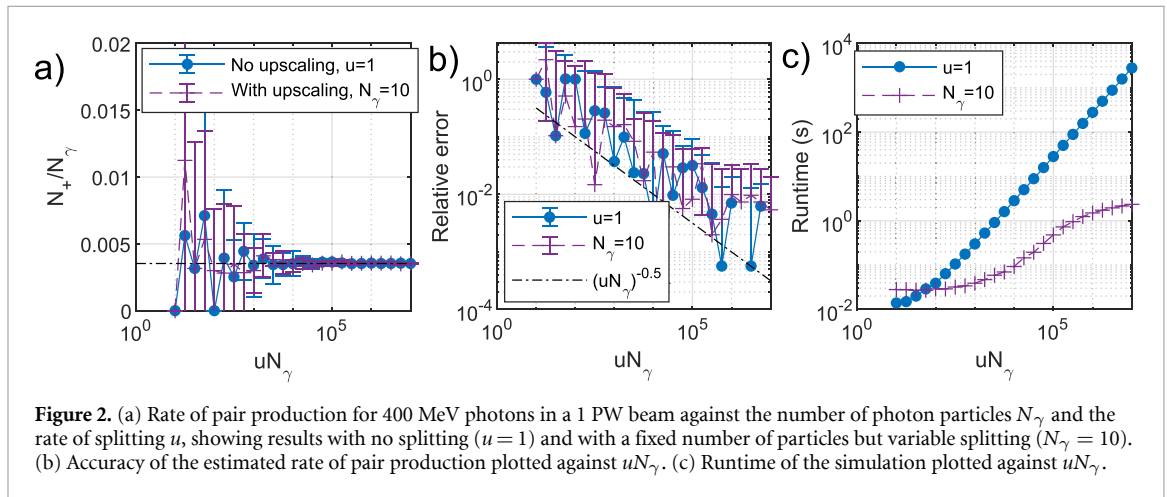
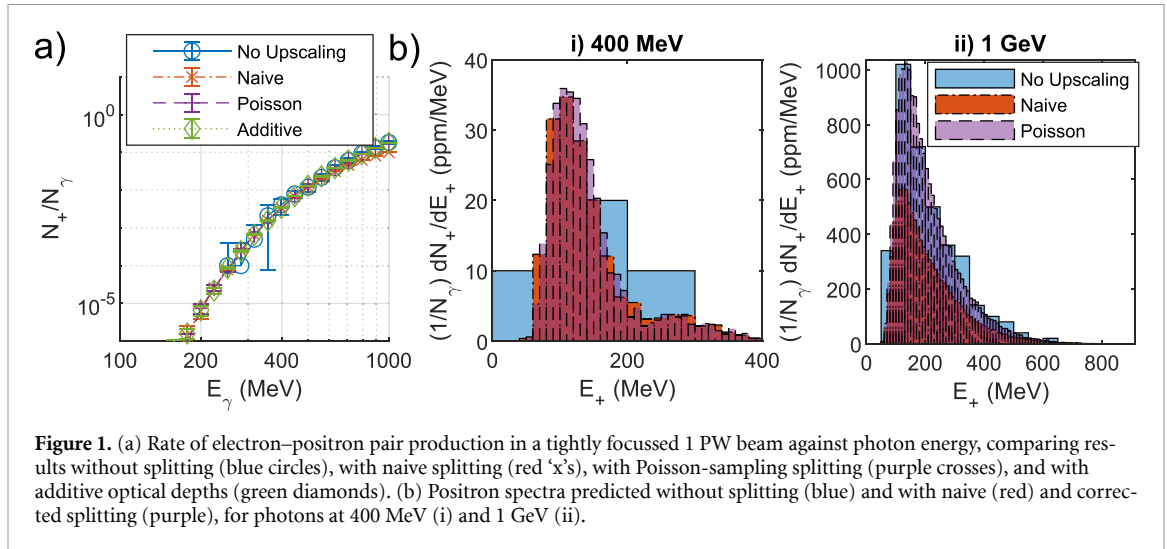
II—Adding the new random optical depth onto the previous value, including the ‘overshoot’ by which the optical depth fell below zero. Another emission occurs when this additive optical depth falls below 0 on a subsequent timestep. As the optical depth becomes the sum of all the random numbers $-\sum_i \ln(1 - r_i)$, minus the total accumulated event path length $\sum_n \lambda t_n$, this maintains both the history of the particle and the independence of the emissions, while using only a single random number for each emission event. In this case, the parent particle weight is not updated after each emission, in order to make sure the emissions are still independent. Instead, the particle is killed with a probability of $1/u$, such that at higher splitting rates the majority of particles are unaffected by emissions. This is a natural extension for codes that use optical depths, like EPOCH [22] or Smilei [24]. It makes pair production events independent again while conserving energy on average.

III - ‘Sub-cycling’ the addition of multiple random optical depths onto the previous value to represent multiple emissions. Typically the ‘additive optical depth’ approach will not allow for more than one emission per timestep, so will also fail for high rates of pair production $\lambda \sim 1$, where more than one pair can be produced by a single macroparticle per timestep. However, it is possible to repeat the process multiple times, continuing to compute emissions so long as the optical depth remains below zero. In practice, this repeated optical depth approach is the technique used by many Poisson sampling algorithms when the mean is not much greater than one, and so if λ is around one or less this sub-cycling approach is still more efficient than Poisson sampling. However, sub-cycling still does not update the parent particle weight after each emission and so only conserves energy on the ensemble average.

The four approaches are summarised in table 1. The most expensive and most accurate approach is full Poisson sampling, which allows for multiple emissions per timestep at very high field strengths and conserves energy on every emission, but requires calculating a random number every timestep (more than one random number if there is more than one emission) and is therefore more expensive. It also requires modelling particles with a wide range of different weights ($w - n/u$ and n/u for any integer n), which can add computational challenges. The naive optical depth approach and additive optical depth approaches are equally computationally cheap, but using an additive optical depth extends the range where upscaling is valid to areas where a particle will have more than one emission over the course of the simulation, at the cost of only conserving energy on average. Subcycling extends the range of applicability still further by allowing multiple emissions per timestep, at the cost of more computation only when multiple emissions per timestep are required. From here on we will only use the version of the additive optical depth with sub-cycling included. All of these optical depth approaches have simpler particle weights, either ‘big’ macroparticles with weight w or ‘small’ macroparticles with weight $1/u$. This makes subsequent numerical behaviour simpler, such as by ignoring ‘small’ particles in subsequent collisions.

2.3. Particle splitting simulations

We tested the three techniques for particle splitting in Monte-Carlo simulations of a head-on collision between high energy photons and an intense laser pulse using the QED-Cascade code [25, 26]. A 1



PW pulse with 25 J in 25 fs was simulated, focussed to a spot waist of $w_0 = 2 \mu\text{m}$ and a peak intensity of $2.2 \times 10^{22} \text{ W cm}^{-2}$ or $a_0 \approx 100$. A 1000 photon particles, each with an initial weight of $w = 1$, were collided with the laser pulse at an angle of 15° , intersecting the peak laser intensity at best focus. The photon beam was a mono-energetic and divergence-free point source with energies varying from 100 MeV up to 1 GeV. 10 simulations were run at each photon energy, with the results shown in figure 1. This shows both the number of pairs produced and the change in the predicted energy spectrum of the positrons.

At photon energies below 300 MeV, the rate of pair production is sufficiently small (below 10^{-4} positrons per photon) that the original simulations without splitting do not produce any positron–electron pairs, leaving the simulations unable to capture the Breit–Wheeler process. Even above this, very small numbers of pairs are produced, giving significant variation from shot noise, as shown by large error bars at each point. At 400 MeV, only four positron macroparticles were produced, giving large shot noise and making it hard to resolve the positron energy distribution. For convergence testing at these low energies we require simulations with more photon macroparticles, described later (see figure 2).

A naive attempt to increase the number of positrons produced simply increased the rate of pair production and reduced the weight of the pair macroparticles produced. This can capture the Breit–Wheeler process at much lower pair production rates, down to 10^{-7} pairs per photon in this case. For photon energies between 300 and 600 MeV this naive splitting gives results that agree well with the original simulation but with many more positron macroparticles and hence much lower shot noise (for instance producing 3400 macroparticles at 400 MeV). However, above 600 MeV, the naive splitting underestimates the rate of pair production by a factor of around two. This is shown most clearly in the predicted positron energy distribution at 1 GeV, demonstrating that the naive splitting is not reliable.

Using both of the corrected splitting methods, returning to sampling the Poisson distribution and adding new optical depths onto the cumulative total, the simulations are able to accurately resolve the pair production rate both at photon energies as low as 150 MeV, and up to 1 GeV. Here, the shot noise is lower and the larger number of positron macroparticles (100 000 at 1 GeV) enables fine resolution of the pair production rates and the positron energy distributions. Both corrected splitting methods are therefore effective for simulating Breit–Wheeler pair production across a wider range of photon energies, even with a relatively small number of photon macroparticles. For the rest of the paper, we focus on the Poisson approach, which is closer to the true physical picture at the cost of a higher runtime.

We can compare the accuracy of simulations with and without splitting by running simulations at a range of splitting rates u and number of photons N_γ for a fixed photon energy of 400 MeV, again colliding with a laser pulse of 25 J in 25 fs. Increasing either u or N_γ should increase the number of positron macroparticles produced and hence improve the estimate of the underlying pair production rate, but will also increase the runtime.

Initially, we run simulations with no splitting ($u = 1$), but with the number of photon macroparticles varying from just $N_\gamma = 10$ up to 10 million; the results are shown in figure 2. As the number of photon macroparticles increases, the rate of pair production predicted by the simulations converges to a consistent value, at $(3.53 \pm 0.02) \times 10^{-3}$ positrons per photon. Both the error on the rate of pair production and the variation across simulations drop with the increasing number of photon macroparticles, limited by the shot noise as expected ($1/\sqrt{N_+} \propto 1/\sqrt{uN_\gamma}$) until the relative error is below 1%. The runtime, however, increases to over 2800 s per simulation, meaning the set of 10 simulations takes almost 8 h. The increase in accuracy therefore comes at a significant cost.

Keeping the number of macroparticles fixed at $N_\gamma = 10$ and increasing splitting, on the other hand, increases the accuracy of the prediction while having much less effect on the runtime. The rate of pair production converges to the same value of $(3.55 \pm 0.05) \times 10^{-3}$ positrons per photon and the accuracy improves at the same rate, $\propto 1/\sqrt{uN_\gamma}$. The runtime, however, only reaches 2.3 s per simulation, three orders of magnitude faster for the same accuracy, meaning that using particle splitting is an efficient way of accurately simulating Breit–Wheeler pair production.

3. Monte-Carlo simulations of realistic experiments

3.1. Simulating shot-to-shot jitter

Running many simulations quickly is particularly useful when optimising parameters for real laser experiments, where there is significant random ‘jitter’ in experimental parameters from shot to shot that reduce the peak quantum parameter of the collision. Even if the energy of the electron beam and the colliding laser are very stable, changes in the laser pointing often cause the electron beam position to move by tens of microns between shots. Similarly, even if the synchronisation of the two laser beams is controlled to a few parts per million, tens of femtosecond shifts in timing are common. When the focus of the colliding beam is only microns across and a few tens of femtoseconds in duration, these effects reduce the measured rate of pair production by orders of magnitude relative to the best-case head-on collision. Simulations using the ideal parameters are often unrealistically optimistic about what measurements are possible.

Now we have demonstrated that particle splitting gives significant reductions in runtime for the same level of accuracy, we are able to run many consecutive Monte-Carlo simulations to model the effect of jitter and capture pair production even at low laser intensity. By modelling hundreds or thousands of consecutive laser shots with the experimental parameters randomly drawn from the known distributions, we can better simulate a realistic laser experiment and estimate the total accumulated rate of pair production. Note, for benchmarking purposes we conduct most of these simulations with mono-energetic electrons, with zero energy spread and no jitter in the electron energy.

An example of such a simulated experiment is shown in figure 3. An initial electron beam with an energy of 7.5 GeV collides with an intense laser pulse, again with 25 J in 25 fs focussed to a $2 \mu\text{m}$ spot. The colliding laser parameters are chosen to reflect the designs of the ZEUS and CoReLS laser facilities [2, 27] (see also the APOLLON laser [28]), while these high electron energies have recently been achieved in guided laser wakefield acceleration experiments [29–31]. Now the prior distributions of laser pointing Δx and Δy (both $\sim \mathcal{N}(0, 10 \mu\text{m})$) and timing jitter $\Delta t \sim \mathcal{N}(0, 25\text{fs})$ lead to imperfect collisions with a distribution of χ_e and χ_γ . The posterior distribution for the number of pairs produced

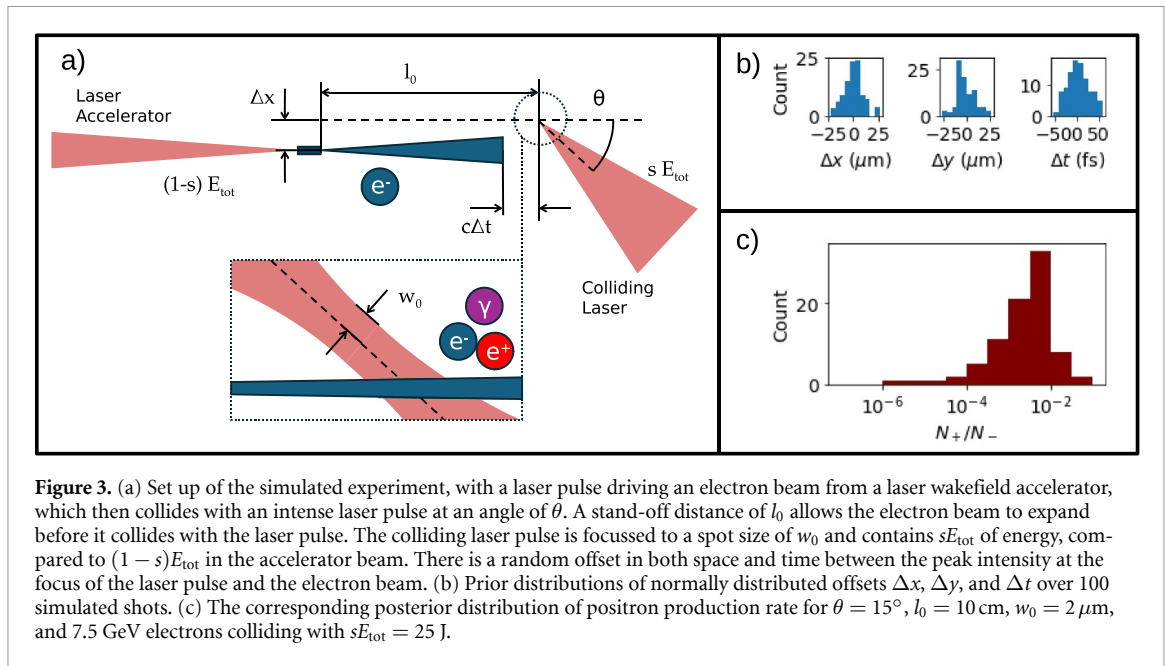


Figure 3. (a) Set up of the simulated experiment, with a laser pulse driving an electron beam from a laser wakefield accelerator, which then collides with an intense laser pulse at an angle of θ . A stand-off distance of l_0 allows the electron beam to expand before it collides with the laser pulse. The colliding laser pulse is focussed to a spot size of w_0 and contains sE_{tot} of energy, compared to $(1-s)E_{\text{tot}}$ in the accelerator beam. There is a random offset in both space and time between the peak intensity at the focus of the laser pulse and the electron beam. (b) Prior distributions of normally distributed offsets Δx , Δy , and Δt over 100 simulated shots. (c) The corresponding posterior distribution of positron production rate for $\theta = 15^\circ$, $l_0 = 10$ cm, $w_0 = 2 \mu\text{m}$, and 7.5 GeV electrons colliding with $sE_{\text{tot}} = 25$ J.

N_+/N_- includes both an error from shot noise and a much more substantial variation from changes in jitter. This simulated experiment took 500 s to run on a single core, giving an average rate of pair production of $N_+/N_- \approx 0.013 \pm 0.001$. This is two orders of magnitude lower than for the ideal ‘head-on’ collision with $\Delta x = \Delta y = \Delta t = 0$, which gives a positron rate of $N_+/N_- \approx 1.2$.

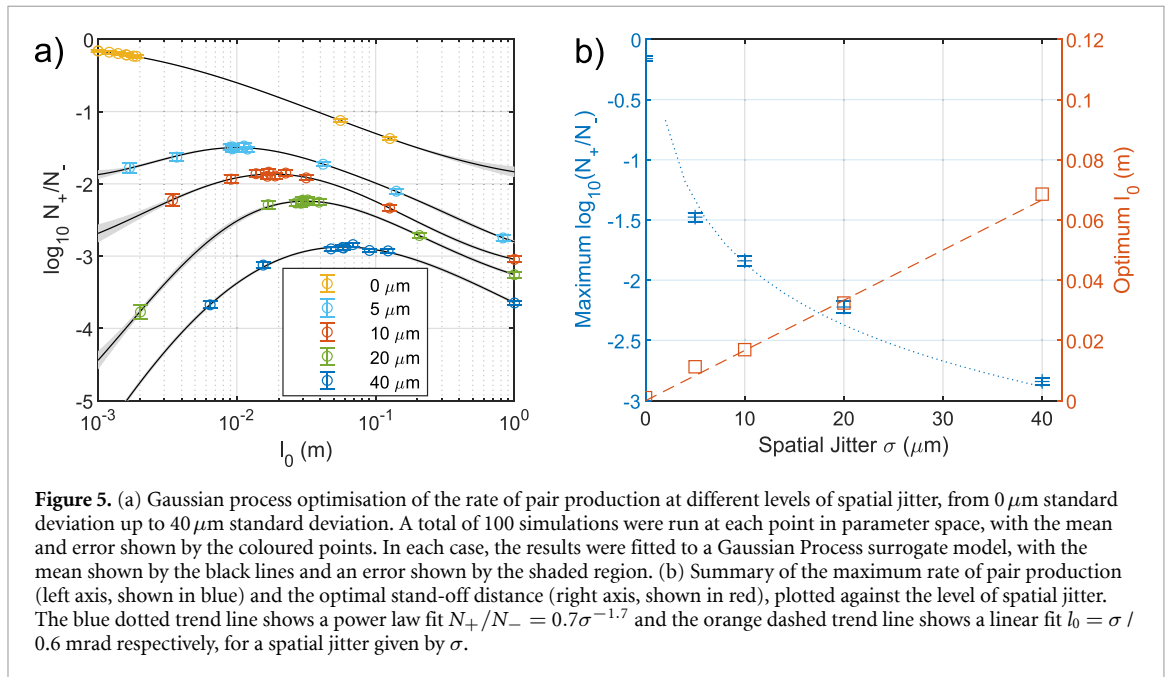
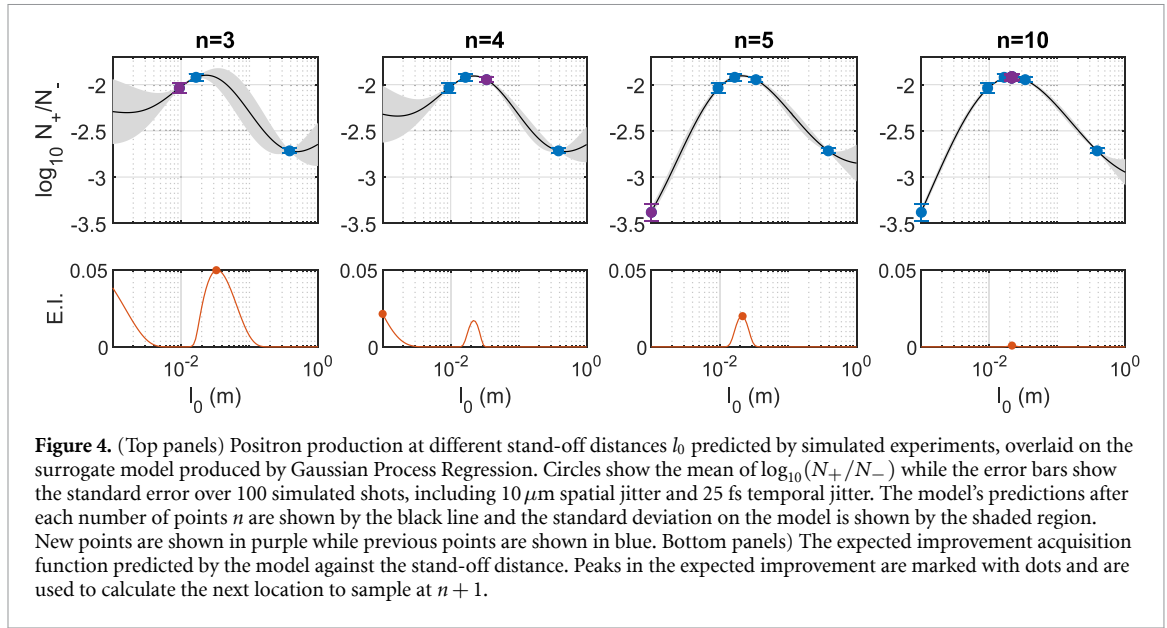
3.2. Optimising experimental conditions

We can now use simulated experiments to optimise the rate of pair production in a laser collider set-up using Gaussian Process Regression. By running many Monte-Carlo simulations at each set of nominal experimental parameters—with variations around those conditions randomly drawn from given distributions—we can determine what parameters to aim for when trying to measure Breit–Wheeler pair production in a real laser experiment. This is a noisy function over a multi-dimensional parameter space, with significant errors from both jitter in the experimental parameters and stochastic noise in the rate of pair production. A Bayesian optimisation approach using Gaussian Process Regression is therefore well suited to rapidly exploring the parameter space and then focussing measurements on the region with the best chances of producing high numbers of pairs.

A Bayesian optimisation algorithm attempts to efficiently maximise a given output from a noisy function, in this case the number of pairs produced in a Monte-Carlo simulation. It does so by building a surrogate model with an associated error, which acts as a proxy for the underlying real function while being much faster to evaluate. In this case, for the surrogate model we use a Gaussian process, which is a stochastic process with a normally distributed error. This continuously interpolates between measurements in the multi-dimensional space so as to ensure that the surrogate model is consistent with the data points and the specified error.

Every time the algorithm makes a new measurement it uses the surrogate model to calculate an acquisition function, takes a measurement where the acquisition function is highest, and uses the results of these measurements to update the surrogate model with the new information. The acquisition function we used is the Expected Improvement, which integrates over the error distribution of the surrogate model to find the place in parameter space where the expected increase from the current maximum to the surrogate model output is highest. This includes locations both where the surrogate model is high, and also where the error on the surrogate model is high, reflecting our lack of confidence in the surrogate to accurately represent the underlying function. Over many measurements, the surrogate model becomes an increasingly accurate proxy for the underlying function, the error on the surrogate model drops, and the acquisition function predicts smaller improvements.

To begin with, we optimised a simulated experiment with just a single variable, the stand-off distance l_0 , for a fixed electron beam divergence of 1 mrad FWHM. The larger the stand-off distance, the larger the electron beam at the collision point, presenting a bigger target to hit. However, a larger electron beam also means a lower electron density and fewer electrons participating in the collision when one occurs. These simulations used 7.5 GeV electrons colliding with $E_{\text{coll}} = 25$ J of laser energy at an



angle of $\theta = 15^\circ$, where the colliding laser pulse was focussed to $w_0 = 2\ \mu\text{m}$ and had a duration of 25 fs FWHM, giving a peak power of around 1 PW.

The results from the GPR algorithm are shown in the top panels of figure 4 as more simulated experiments are run at different values of the stand-off distance l_0 . Initially, after just $n = 3$ initial random points, the model carries a high uncertainty across most of the domain. As more data points are taken this uncertainty gets smaller until by $n = 5$ the underlying function near the maximum and the position of the maximum are known accurately, while uncertainty remains elsewhere in the domain. Further points up to $n = 10$ continue to reduce the error near the maximum. After each simulation, the next set of simulated experiments are run where the expected improvement (shown by the bottom panels) is high, narrowing down the possibilities. As more points are sampled and the confidence of the model increases, the expected improvement falls. With the ‘explore’ parameter set to just 0.2, the algorithm increases the expected improvement in the region close to the current maximum rather than in other regions of the domain. Ultimately, this gives a maximum of 0.015 ± 0.001 positrons per initial electron at a stand-off distance of $l_0 = 17\ \text{mm}$.

We can also look at the impact of changing the other experimental conditions, such as the jitter, with examples shown in figure 5. With no spatial jitter the maximum rate of pair production is

0.69 ± 0.03 pairs produced per initial electron, with zero stand-off-distance, reflecting the high electron energy and laser intensity. However, increasing the spatial jitter causes a reduction in the average positron production rate by orders of magnitude, as the likelihood of an electron hitting the most intense part of the laser at a given stand-off distance reduces as σ^{-2} for a spatial jitter given by σ . This effect can be ameliorated by increasing the stand-off distance and hence the size of the electron beam at the collision point, with the optimal stand-off distance increasing to 69 mm at $\sigma = 40 \mu\text{m}$. At these large stand-off distances, the larger electron beam is easier to hit and the increase in the number of shots that hit the electron beam outweighs the disadvantage of only a small fraction of the electron beam interacting with the laser pulse on each shot. At higher spatial jitters the optimum stand-off distance increases linearly such that the optimal size of the electron beam should be comparable to the spatial jitter.

3.3. Optimum parameters for non-linear Breit–Wheeler pair production

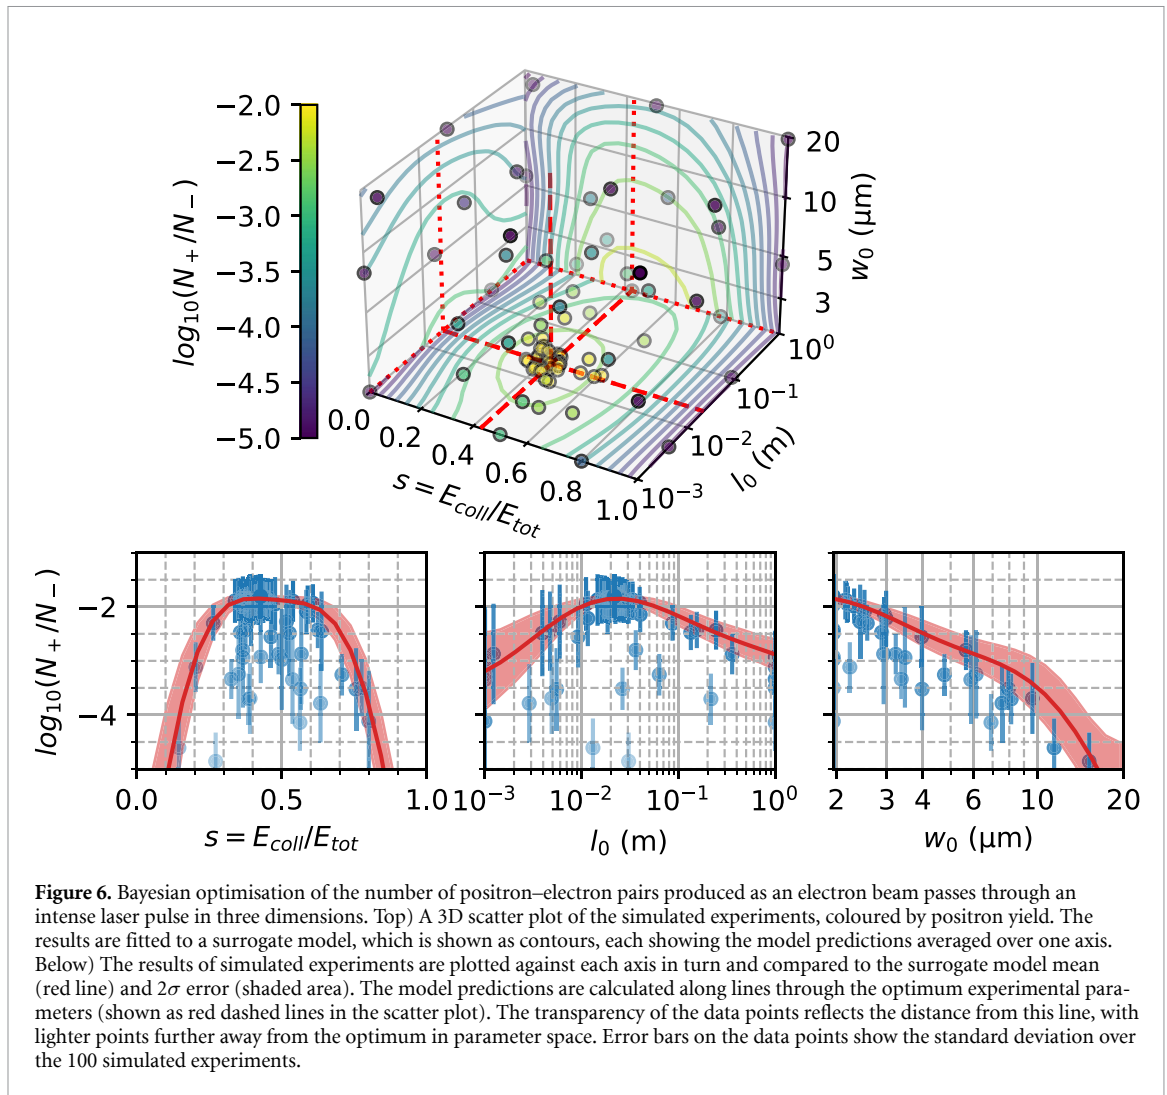
Now we have demonstrated a one-dimensional optimisation, we will use Gaussian Process Regression to maximise the number of positrons produced across several parameters in simulated experiments. To find the optimum parameters, we consider changes in the energy split $s = E_{\text{coll}}/E_{\text{tot}}$, for $E_{\text{tot}} = 100$ J. The resulting electron energy is assumed to be linear with laser energy as $E_e = 10(1-s)$ GeV. This scaling reflects all-optical inverse Compton scattering experiments over the last 13 years at a range of facilities [32–36]. We also vary the laser waist w_0 between $2 \mu\text{m}$ and $20 \mu\text{m}$, and the stand-off distance l_0 between 1 mm and 1 m. A smaller laser spot and shorter duration means the laser reaches a higher intensity, enabling larger quantum parameters χ_e and χ_γ . However, it also provides a smaller target such that when the spatial or temporal jitter is large, the electron beam is more likely to miss the laser pulse completely and encounter only a very low quantum parameter. Similarly, a smaller stand-off distance means the electron bunch is small at the collision point, with almost all electrons encountering a high laser intensity, but it increases the probability of the electron beam missing the colliding laser entirely. Larger stand-off distances increase the probability of a collision but reduce the number of electrons taking part in the interaction. We use a particle splitting rate of $u = 1000$ to better resolve the positron production rate in lower field regions of the parameter space and reduce the error from shot noise.

Until now, we have performed the simulations for monoenergetic photons (section 2) and monoenergetic electrons (sections 3.1 and 3.2). In practice, however, LWFA electron beams typically have high variability shot-to-shot and significant energy spread. For these simulations we model energy spread by sampling electrons from a Gaussian distribution with a standard deviation 10% of the electron mean energy. Shot-to-shot variation in the electron mean energy itself is modelled by sampling from a Gaussian distribution with a mean of $10(1-s)$ GeV and a standard deviation of $(1-s)$ GeV, where $s = E_{\text{coll}}/E_{\text{tot}}$ as before. This reflects approximately 10% shot-to-shot variation in electron energy and 10% energy spread, as in e.g. references [31, 37]. The electron beam divergence remains fixed at 1 mrad. These parameters are representative of all-optical collision experiments, but are not tied to a particular setup or facility.

Figure 6 shows the results of the simulations, and the Bayesian optimisation as a series of projections through the data. The maximum number of positrons obtained in these simulated experiments with $10 \mu\text{m}$ spatial jitter is $(1.57 \pm 0.18) \times 10^{-2}$, or around 2 pairs per 100 electrons. This optimum is obtained at an energy split of $s = 0.43$, giving a mean electron energy of 5.7 GeV colliding with 1.7 PW of laser power. As in the previous section, the positron yield is maximised at a stand-off distance where the electron beam size is comparable to the $10 \mu\text{m}$ spatial jitter, here at $l_0 = 22$ mm. The minimum laser waist of $2 \mu\text{m}$ gives a peak field strength of $a_0 = 132$ and yields the largest number of positrons, with the rate of pair production falling four orders of magnitude with a factor of 10 increase in w_0 . This highlights the extremely strong dependence of Breit–Wheeler pair production on the quantum parameter, which here would nominally reach up to $\chi_e \approx 9$ for an ideal collision with no radiation reaction. These simulated experiments demonstrate that pair production requires the highest possible laser intensity at focus, even if it means accepting a relatively large stand-off distance, with few electrons reaching the peak field strength.

3.4. Optimum parameters for measuring inverse Compton scattering

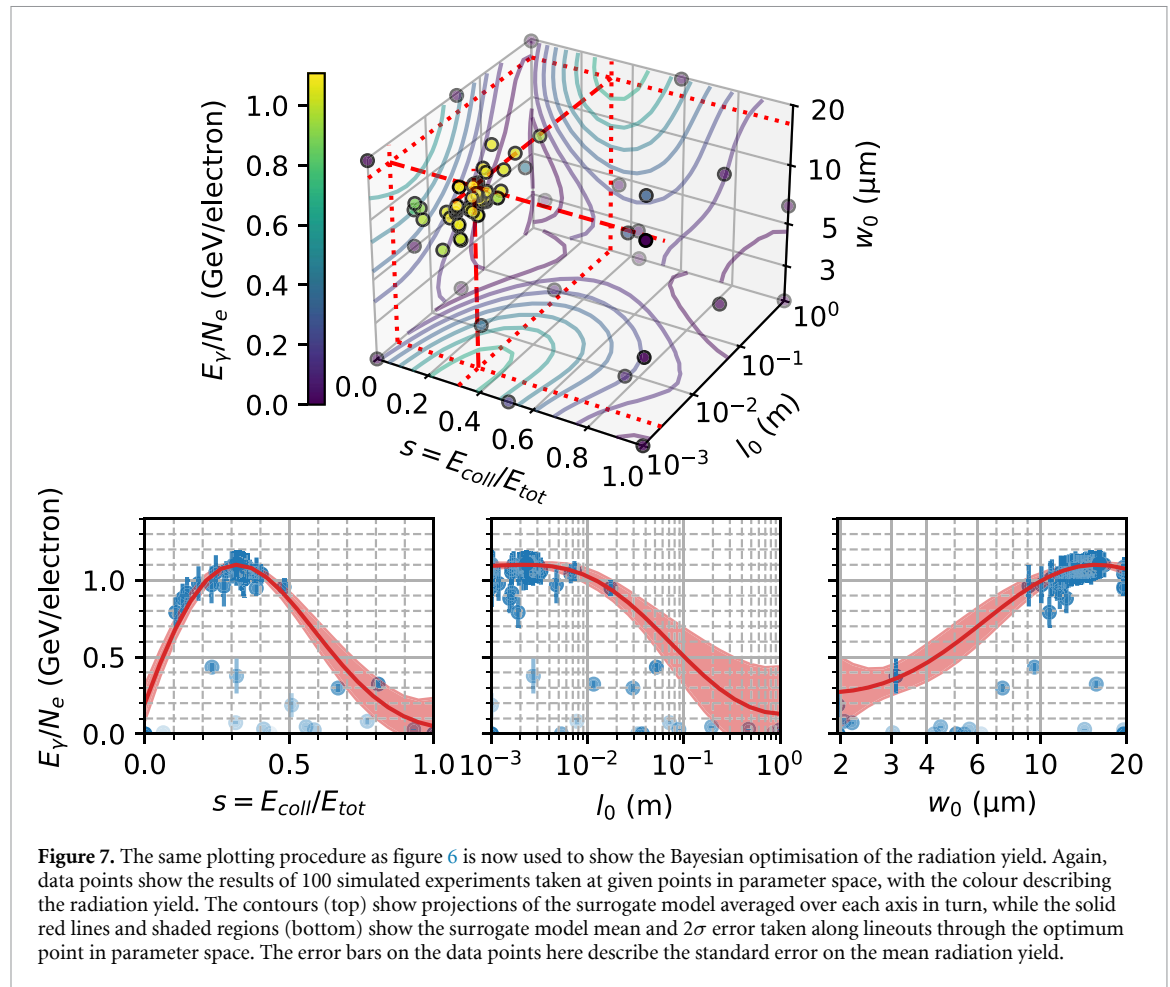
We can also use simulated experiments to maximise the synchrotron radiation emitted by an electron beam as it collides with an intense laser pulse. This radiation is emitted at energies up to the 10 s–100 s of MeV scale and can contain a significant fraction of the electron energy. The high-energy gamma rays produced are a key result of the many-photon interactions in the strong field regime and can lead to



the production of positron–electron pairs. Generally the energy in the radiation is increased by higher electron energies and strong field strengths (through the quantum parameter χ_e), and a longer interaction time. With the significant jitter of these simulated experiments it is also important to maximise the number of particles involved in the interaction. We vary the same three parameters the energy split $s = E_{\text{coll}}/E_{\text{tot}}$, the stand-off distance l_0 , and the laser waist w_0 .

Figure 7 shows the result of the three-dimensional maximisation of the radiated energy. Points selected for simulated experiments are rapidly focussed in on the most promising region, giving a more accurate surrogate model with lower error bars near the maximum. This allows us to quickly determine the best parameters at which to run experiments to produce the maximum amount of radiation from an electron beam colliding with a laser pulse. In this situation, with 100 J of laser energy available, the optimum is found with 68 J used to accelerate the electron beam to 6.8 GeV, and 32 J used to collide with it. This matches the naive expectation of maximising the peak quantum parameter $\chi_e \propto \gamma\sqrt{I_0} \propto (1-s)\sqrt{s}$.

However, the laser waist which maximises the radiated energy is now around $16\ \mu\text{m}$, much larger than the tightest possible spot and resulting in a peak intensity of $a_0 \approx 14$, much lower than the maximum achievable $a_0 = 200$. This increases the number of successful collisions by greatly increasing the volume around focus. When the laser waist is this large, the electron beam can be small and still collide effectively, such that the maximum energy radiated is again located at the minimum stand-off distance of $l_0 = 1\ \text{mm}$.



4. Conclusion

In this work, we have demonstrated steps to efficiently optimise experimental conditions for all-optical Breit–Wheeler pair production using Monte-Carlo simulated experiments. First, we introduce two new algorithms for particle splitting which enables us to quickly and accurately sample the probability of producing a positron–electron pair many times within a single macroparticle. In contrast to the naive approach using a single optical depth, these techniques correctly predict the rate of pair production over at least an order of magnitude of variation in the quantum parameter χ_γ , while giving runtimes up to a thousand times faster than by increasing the number of macroparticles. Secondly, this makes it practical to run many Monte-Carlo simulations at a single set of parameters in order to account for experimental jitter in laser pointing and timing. Given prior distributions for the shot-to-shot jitter, these simulated experiments predict a posterior distribution of rates of pair production, which gives the practical likelihood of measuring positron–electron pairs in experiments.

This approach makes it possible to efficiently optimise experimental parameters, maximising the likelihood of measuring Breit–Wheeler pair production in laser-driven experiments. Gaussian Process regression was used to build surrogate models of the rate of pair production and synchrotron radiation produced in simulated collisions between an LWFA electron beam and an intense laser pulse. These showed that laser pointing jitter can be ameliorated by increasing the size of the laser focal spot or by allowing the electron beam to expand after leaving the wakefield accelerator, with a stand-off distance between the accelerator and the collision point. Synchrotron radiation, which requires a lower laser intensity than pair production, was maximised with a larger focal spot ($\sim 15 \mu\text{m}$), and no stand-off distance. In contrast, a stand-off distance of a few centimetres maximised the rate of pair production by increasing the size of the electron beam to tens of microns, comparable to the scale of the pointing jitter. This increases the number of electrons which collide with the laser focus, a focus which should be as small as possible (a few microns) in order to reach the highest possible laser intensity and quantum parameter. Together, these parameters enable near-term laser facilities to produce pairs at a rate of around one pair per one hundred electrons, even with realistic timing and pointing jitter.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: https://github.com/ChrisArran/gp_optimise [38].

Funding

This work was funded by EPSRC Grant EP/V049461/1 and in part by the Air Force Office of Scientific Research (FA9550-24-1-0053). The Viking cluster was used during this project, which is a high performance compute facility provided by the University of York. We are grateful for computational support from the University of York IT Services and the Research IT team.

ORCID iDs

Christopher Arran  0000-0002-8644-8118

Stuart Morris  0000-0002-2218-0706

Christopher P Ridgers  0000-0002-4078-0887

References

- [1] Radier C *et al* 2022 *High Power Laser Sci. Eng.* **10** e21
- [2] Yoon J W, Kim Y G, Choi I W, Sung J H, Lee H W, Lee S K and Nam C H 2021 *Optica* **8** 630–5
- [3] Bell A R and Kirk J G 2008 *Phys. Rev. Lett.* **101** 200403
- [4] Bulanov S S, Esirkepov T Z, Thomas A G R, Koga J K and Bulanov S V 2010 *Phys. Rev. Lett.* **105** 220407
- [5] Ridgers C P, Brady C S, Ducloux R, Kirk J G, Bennett K, Arber T D, Robinson A P L and Bell A R 2012 *Phys. Rev. Lett.* **108** 165006
- [6] Schwinger J 1951 *Phys. Rev.* **82** 664–79
- [7] Breit G and Wheeler J A 1934 *Phys. Rev.* **46** 1087–91
- [8] Burke D L *et al* 1997 *Phys. Rev. Lett.* **79** 1626–9
- [9] Bula C *et al* 1996 *Phys. Rev. Lett.* **76** 3116–9
- [10] Esberg J *et al* 2010 *Phys. Rev. D* **82** 072002
- [11] Nielsen C F, Holtzapfle R, Lund M M, Surrow J H, Sørensen A H, Sørensen M B and Uggerhøj U I 2023 *Phys. Rev. Lett.* **130** 071601
- [12] Nees J *et al* 2020 ZEUS: a national science foundation mid-scale facility for laser-driven science in the QED regime 2020 *Conf. Lasers and Electro-Optics (CLEO)* pp 1–2
- [13] Piazza A D, Willingale L and Zuegel J D 2022 Multi-petawatt physics prioritization (MP3) *Workshop Report* (<https://doi.org/10.48550/arXiv.2211.13187>)
- [14] Timokhin A N 2010 *Mon. Not. R. Astron. Soc.* **408** 2092–114
- [15] Vranic M, Grismayer T, Martins J L, Fonseca R A and Silva L O 2015 *Comput. Phys. Commun.* **191** 65–73
- [16] Dong Q, Wang B, Duan X, Yan L, Liu K, Luo W, Shao F and Yu T 2024 *Comput. Phys. Commun.* **294** 108913
- [17] Kawrakow I, Rogers D W O and Walters B R B 2004 *Med. Phys.* **31** 2883–98
- [18] Ivanchenko V N *et al* (EDP Sciences) 03101
- [19] Ramos-Méndez J, Schuermann J, Incerti S, Paganetti H, Schulte R and Faddegon B 2017 *Phys. Med. Biol.* **62** 5908–25
- [20] Smets R, Aunai N, Ciardi A, Drouin M, Campos-Pinto M and Deegan P 2021 *Comput. Phys. Commun.* **261** 107666
- [21] Arber T D *et al* 2015 *Plasma Phys. Control. Fusion* **57** 113001
- [22] Ridgers C, Kirk J, Ducloux R, Blackburn T, Brady C, Bennett K, Arber T and Bell A 2014 *J. Comput. Phys.* **260** 273–85
- [23] Blackburn T G, King B and Tang S 2023 *Phys. Plasmas* **30** 093903
- [24] Derouillat J *et al* 2018 *Comput. Phys. Commun.* **222** 351–73
- [25] Watt R 2021 Monte Carlo Modelling of QED Interactions in Laser-Plasma Experiments *PhD Thesis* Imperial College London (<https://doi.org/10.25560/91984>)
- [26] Watt R 2021 QED-Cascade (available at: <https://github.com/robbiewatt1/QED-Cascade>) (Accessed 09 April 2026)
- [27] Willingale L *et al* 2023 Status of the ZEUS laser user facility 2023 *Conf. Lasers and Electro-Optics (CLEO)* pp 1–2
- [28] Zou J *et al* 2015 *High Power Laser Sci. Eng.* **3** e2
- [29] Gonsalves A J *et al* 2019 *Phys. Rev. Lett.* **122** 084801
- [30] Miao B *et al* 2022 *Phys. Rev. X* **12** 031038
- [31] Picksley A *et al* 2024 *Phys. Rev. Lett.* **133** 255001
- [32] Tsai H-E *et al* 2016 *Phys. Plasmas* **22** 023106
- [33] Ta Phuoc K, Corde S, Thauray C, Malka V, Tafzi A, Goddet J P, Shah R C, Sebban S and Rousse A 2012 *Nat. Photon.* **6** 308–11
- [34] Chen S *et al* 2013 *Phys. Rev. Lett.* **110** 155003
- [35] Poder K *et al* 2018 *Phys. Rev. X* **8** 031004
- [36] Mirzaie M *et al* 2024 *Nat. Photon.* **18** 1212–7
- [37] v Grafenstein K *et al* 2023 *Sci. Rep.* **13** 11680
- [38] Arran C 2026 *gp_optimise* (available at: https://github.com/ChrisArran/gp_optimise)