



Optimal parameters for radiation reaction experiments

Downloaded from: <https://research.chalmers.se>, 2025-12-10 00:27 UTC

Citation for the original published paper (version of record):

Arran, C., Cole, J., Gerstmayr, E. et al (2019). Optimal parameters for radiation reaction experiments. Plasma Physics and Controlled Fusion, 61(7). <http://dx.doi.org/10.1088/1361-6587/ab20f6>

N.B. When citing this work, cite the original published paper.

Optimal parameters for radiation reaction experiments

C Arran¹ , J M Cole², E Gerstmayr² , T G Blackburn³ ,
S P D Mangles² and C P Ridgers¹

¹ York Plasma Institute, University of York, York, United Kingdom

² The John Adams Institute for Accelerator Science, Imperial College London, London, United Kingdom

³ Department of Physics, Chalmers University of Technology, Gothenburg, Sweden

E-mail: christopher.arran@york.ac.uk

Received 25 January 2019, revised 26 April 2019

Accepted for publication 10 May 2019

Published 11 June 2019



Abstract

As new laser facilities are developed with intensities on the scale of 10^{22} – 10^{24} W cm⁻², it becomes ever more important to understand the effect of strong field quantum electrodynamic processes, such as quantum radiation reaction, which will play a dominant role in laser-plasma interactions at these intensities. Recent all-optical experiments, where GeV electrons from a laser wakefield accelerator encountered a counter-propagating laser pulse with $a_0 > 10$, have produced evidence of radiation reaction, but have not conclusively identified quantum effects nor their most suitable theoretical description. Here we show the number of collisions and the conditions required to accomplish this, based on a simulation campaign of radiation reaction experiments under realistic conditions. We conclude that while the critical energy of the photon spectrum distinguishes classical and quantum-corrected models, a better means of distinguishing the stochastic and deterministic quantum models is the change in the electron energy spread. This is robust against shot-to-shot fluctuations and the necessary laser intensity and electron beam energies are already available. For example, we show that so long as the electron energy spread is below 25%, collisions at $a_0 = 10$ with electron energies of 500 MeV could differentiate between different quantum models in under 30 shots, even with shot-to-shot variations at the 50% level.

Keywords: radiation reaction, Monte-Carlo simulations, high field physics, laser-plasma interactions

(Some figures may appear in colour only in the online journal)

1. Introduction

Experiments using new ultra-high-intensity multi-petawatt laser facilities such as Apollon 10 PW [1] and ELI [2, 3] will require a thorough experimental understanding of non-classical behaviour in laser-plasma interactions. In experiments reaching laser intensities of 10^{22} – 10^{24} W cm⁻² the effect of strong field QED processes starts to strongly modify the laser-plasma interaction [4–6], and better understanding the

fundamental physical processes at work will be crucial. One of these processes, the radiation produced by charged particles when moving in an electro-magnetic field, and the subsequent recoil experienced by the particles, is particularly relevant to studies of inverse Compton scattering [7, 8] and laser absorption in solid target interactions [9–12], both of which are key targets of the ELI-NP facility.

Two recent experiments [13, 14] have aimed to study the effect of radiation reaction in isolation, using existing petawatt laser facilities with peak intensities of $I \sim 10^{21}$ W cm⁻². In these all-optical set ups, highly energetic electrons ($\gamma = 1000$ – 2000) produced by a laser wakefield accelerator [15–18] collide with a counter-propagating high intensity laser pulse as shown in the schematic in figure 1. In the rest



Original content from this work may be used under the terms of the Creative Commons Attribution 3.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.




Figure 1. Schematic of an all-optical radiation reaction experiment. An intense ultra-short laser pulse, the LWFA drive beam, is incident upon a gas target, producing a high energy electron bunch. A second intense laser pulse, the colliding beam, is brought to a tight focus just outside the gas target, interacting with the electron bunch and producing a beam of high energy light.

frame of the electron both the frequency and the intensity of the radiation are dramatically increased, bringing the electric field experienced by the electron to the scale of $E'_L \approx \gamma E_L \sim 10^{17} \text{ V m}^{-1}$, comparable to the Schwinger limit $E_s = 1.32 \times 10^{18} \text{ V m}^{-1}$ [19], as described by the dimensionless and Lorentz invariant parameter $\chi_e = E'_L/E_s$. At this point the predictions from quantum and classical models of radiation reaction strongly diverge; whereas using the classical synchrotron spectrum requires the production of photons with energies $\varepsilon_\gamma > \varepsilon_e$, the quantum model limits the energy of photons so as to conserve energy, significantly reducing the synchrotron power at high field strengths, as described in [20–23]. Both of the recent experiments demonstrate significantly better agreement between their measurements and quantum, non-perturbative, models than with classical models such as described by Landau and Lifshitz [24].

However, the limited number of events measured in the experiments has left significant uncertainty [25], with Poder *et al* [14] concluding a slightly better agreement with a semi-classical model, while the measurements made by Cole *et al* [13] were not able to distinguish between the semi-classical and stochastic models. In the semi-classical description, both the rate of radiation emission and the subsequent change in electron energy are adjusted to match the quantum model, but the emission remains a continuous process, with the recoil a frictional force that leads to cooler electrons with a narrower energy distribution. In the quantum picture, on the other hand, emission is a quantised, stochastic event; some electrons travel much further through the laser pulse before emitting a photon while others emit many, leading to substantial broadening of the electron energy distribution [26–28]. In modelling stochastic emission events, we assume that photon emission is sufficiently fast that the laser field is constant throughout the process, in the so-called constant-cross-field approximation. This is accurate when the coherence time of emission is much less than the laser period, which generally gives a condition $\tau_{\text{COH}} \sim mc/eE_L \ll 1/\omega$, for a laser frequency ω [29], or in terms of the normalised vector potential of the laser pulse, $a_0 = eE_L/m_e\omega c \gg 1$. Even if this condition is met, however, the constant-crossed-field approximation—and both the quantum and semi-classical models—breaks down when the energy of the emitted photon energy is

very low [30, 31], although these photons do not contribute significantly to the recoil [32].

Experiments currently underway seek to resolve the seeming disparity between the two experiments to date and to determine which, if either, model is most appropriate for high intensity laser experiments; this paper attempts to both find the best way of conducting these radiation reaction experiments, and demonstrate the regimes where the choice of model is important. By simulating the experiments under different conditions we place constraints on the parameters required, such as laser intensity and electron energy spread, as well as the accuracy to which these parameters must be controlled. Given different experimental parameters, we estimate the number of measurements required to be confident which model is more appropriate: the stochastic quantum model, or the continuous and deterministic semi-classical model. In doing so, we take account of the shot-to-shot variation of both the energy of electrons from a laser wakefield accelerator and the intensity of the colliding laser pulse.

2. Simulated experiments

In a radiation reaction experiment of the type shown in figure 1 it is important to achieve good overlap in both time and space between the highest intensity region of the laser pulse and the brightest part of the electron beam. However, if the pulse collides with the electron beam close to the LWFA gas target, the electron bunch will be under one micron in diameter. If the collision point is slightly away from the focal plane of the laser, as in [13], it is possible to ensure that the beam profile is much larger than the electron bunch, maximising the overlap in space and the chances of a successful collision. Under these conditions we can reduce the problem to a single dimension, whereas if this is not the case then accurate knowledge of the transverse electron and laser profiles at the collision point are required. Similarly, synchrotron radiation is emitted within a forward-pointing cone with an angle of $1/\gamma$ around the direction of motion of the electron; for an electron bunch with an angular divergence on the scale of $\theta \sim 1 \text{ mrad}$ the total cone angle will be $\approx \theta + a_0/\gamma \sim 10 \text{ mrad}$ in the plane of polarisation of the laser. Under these conditions we will ignore the angular distribution of radiation for energetic electrons. On the short timescale after the creation of the electron beam we can also neglect direct electron–electron interactions such as space-charge, while the electron bunch duration is sufficiently short that we can neglect the interactions of synchrotron photons after they have been emitted. Likewise, as the number of electrons in the bunch is small, we neglect effects of the electron bunch on the laser beam, such as energy loss and refraction. Finally, as we will be considering situations achievable with existing laser facilities, with both $E_L/E_s \ll 1$ and $\chi_e \lesssim 1$, we can neglect pair production. We therefore consider the interaction between each electron and the laser pulse independently, allowing us to reduce a complicated simulation to a sum of many single particle interactions, where each electron has a single initial and final energy but




Figure 2. Contours of (a) the mean final electron energy $\langle \varepsilon_f \rangle$ (b) the final electron energy spread σ_{ε_f} and (c) the critical energy of emitted photons $\varepsilon_{\text{crit}}$, for each of the quantum (solid lines), semi-classical (dashed) and classical (dotted) models. Results are from mono-energetic electron beams in EPOCH simulations. Each line shows the initial electron energy ε_i and the laser a_0 required to obtain the given final state. For the classical and semi-classical models $\sigma_{\varepsilon_f} < 1$ MeV and hence the contours are not visible.

may produce many photons over the course of the interaction with the laser pulse.

First, look-up tables were assembled of the final electron and photon spectra resulting when initially mono-energetic electron beams encountered a laser pulse. One dimensional particle-in-cell simulations with EPOCH [33], employing an extended QED module [22] (see appendix A for details), were conducted for laser peak intensities of $1 \leq a_0 \leq 25$ and electron energies of $100 \leq \varepsilon_i \leq 2$ GeV. These used each in turn of a fully classical, Landau–Lifshitz, radiation reaction model; a semi-classical model with corrected emission rates and powers but with continuous, deterministic emission; and a quantum model of radiation reaction with stochastic emission. By performing 1D simulations we ignored spatial variations in the laser intensity across the electron beam at the collision point. This is valid if the electron beam transverse size is smaller than the laser transverse size at the collision point, as described earlier. The laser pulse had a Gaussian temporal profile with a duration of 40 fs FWHM, chosen to reflect parameters of the recent radiation reaction experiments. Shorter pulse durations would allow experiments to reach higher laser intensities for the same input energy, reaching higher $\chi_e \propto t_{\text{FWHM}}^{-\frac{1}{2}}$ and exploring more non-classical effects⁴. However, in practice, pulse duration is limited by existing laser technology, with shorter pulses typically coming at the expense of laser energy.

For each different reaction model, these look-up tables gave final electron and photon energy distributions $N_{e,f}(\varepsilon_f|\varepsilon_i, a_0)$ and $N_\gamma(\varepsilon_\gamma|\varepsilon_i, a_0)$. The electron energy distributions were fitted to a Gaussian to give functions $\langle \varepsilon_f \rangle(\varepsilon_i, a_0)$ and $\sigma_{\varepsilon_f}(\varepsilon_i, a_0)$, whereas the photon energy distributions were fitted to an expression of the form:

$$N_\gamma \propto \varepsilon_\gamma^{-\frac{2}{3}} \exp\left(-\frac{\varepsilon_\gamma}{\varepsilon_{\text{crit}}}\right), \quad (1)$$

⁴ It is possible to show (see appendix C) that for $\chi_e \ll 1$, reducing the pulse duration for a fixed laser energy only increases the total stochastic broadening, improving the chances of distinguishing between models. For $\chi_e \gtrsim 0.56$, on the other hand, reducing the laser pulse duration is counter-productive.

giving the critical energy $\varepsilon_{\text{crit}}(\varepsilon_i, a_0)$. For examples of the fitted energy distributions, see appendix B. The resulting parameters are plotted in figure 2 using contour plots to show the differences between the three different models.

Firstly, figure 2 demonstrates that the simulations of radiation reaction are working as expected. For a given initial electron energy, increasing the laser intensity reduces the final energy of the electrons, while increasing the initial electron energy leads to the emission of higher energy photons. Both of these correspond to a larger radiation reaction force, with an electron beam losing more power, emitted as photons. However, although at low values of a_0 the photon critical energy increases with laser intensity, for higher a_0 it saturates and for the highest values of ε_i actually begins to decrease. This is because electrons lose so much energy during the radiation reaction process that the peak field experienced by the electrons in their rest frame, γE_L , is actually reduced. In this situation the radiated spectrum comprises a greater number of lower energy photons. Similarly, in the quantum model, electrons experience the greatest stochastic broadening at moderate laser intensities, around $a_0 \approx 10$, while above this the final energy spread is smaller. At moderate laser intensities electrons emit fewer photons on average, with greater variation between electrons due to shot noise.

The look-up tables also allow us to distinguish between the different models for radiation reaction: as laser intensity increases, the classical model predicts much lower final electron energies than the quantum or semi-classical models. Applying the quantum correction—limiting the photon energy to $\varepsilon_\gamma < \varepsilon_e$ —results in significantly higher final electron energies and slightly lower photon energies. The greatest difference in final electron energy occurs at the highest laser intensities and electron energies, whereas the greatest difference in critical energy is centred around $a_0 \approx 10$.

In both the mean final electron energy and photon energy, it is very difficult to see any difference between the predictions from the quantum and semi-classical models. These models contain the same correction to the power radiated, and on average the electrons encounter the same radiation reaction. However, without the effect of stochastic broadening, an initially mono-energetic electron beam




Figure 3. Results from simulated radiation reaction experiments for an initial electron beam with a peak energy of (550 ± 20) MeV and an energy spread of 250 MeV, shown through (a) the three-dimensional joint probability distribution function $f(\langle \epsilon_f \rangle, \epsilon_{\text{crit}}, \sigma_{\epsilon_f})$, and the 2D distribution functions (b) $f_1(\langle \epsilon_f \rangle, \epsilon_{\text{crit}})$ and (c) $f_2(\langle \epsilon_f \rangle, \sigma_{\epsilon_f})$. The 1σ contours are shown, within which 68% of simulated experiments measured these results, assuming each of the classical (yellow), semi-classical (red), and quantum (blue) models. In (a) the intensity of the colliding laser pulse is $a_0 = 11 \pm 3$, whereas for (b) and (c) the intensity was assumed uniformly distributed between $a_0 = 4$ –20, for comparison with figure 9 of [13].

remains mono-energetic in the semi-classical model. In contrast, in the quantum model the electron energy spectrum becomes substantially broader.

Once the look-up tables $\langle \epsilon_f \rangle(\epsilon_i, a_0)$, $\sigma_{\epsilon_f}(\epsilon_i, a_0)$ and $\epsilon_{\text{crit}}(\epsilon_i, a_0)$ were assembled, Monte-Carlo simulated experiments were conducted for more realistic (though still idealised) initial electron energy distributions, which were not mono-energetic, and where each of the laser intensity, the mean electron energy, and the electron energy spread varied shot-to-shot. This allowed us to estimate the underlying three-dimensional probability distribution function $f(\langle \epsilon_f \rangle, \sigma_{\epsilon_f}, \epsilon_{\text{crit}})$ for making measurements of the final mean electron energy $\langle E_f \rangle$, and energy spread σ_{E_f} , and the final photon critical energy ϵ_{crit} . For parameter scans, the two-dimensional probability distribution functions $f_1(\langle \epsilon_f \rangle, \epsilon_{\text{crit}})$ and $f_2(\langle \epsilon_f \rangle, \sigma_{\epsilon_f})$ were calculated instead. The details are described in appendix B.

3. Distinguishing models

First, experimental parameters were chosen to match those in [13], with the laser intensity estimated as $a_0 = 11 \pm 3$ and the electron energy estimated as $\langle \epsilon_i \rangle = (550 \pm 20)$ MeV, with an

energy spread of $\sigma_{\epsilon_i} = 250$ MeV. The three-dimensional distribution functions, shown in figure 3, demonstrate the capability of the simulated experiments. Points in the top right of the image correspond to shots with low a_0 , with high final energies, high energy spread, and lower photon energies. In this regime, the three models predict very similar results. As a_0 increases, the electron beam loses more energy and becomes cooler with a lower energy spread, in the process producing higher energy photons. At the largest values of a_0 , the different rates of radiation reaction and of radiative cooling lead to the three models predicting different results, with the classical model leading to the lowest final energy spread and the highest photon energy, while the quantum model predicts a significantly higher energy spread than either of the two other models. The shot-to-shot variation of $\langle \epsilon_i \rangle$ tends to blur out this trend, broadening the distribution functions and making it more difficult to distinguish between different models.

In order to show this more clearly, and to compare the results with [13], the two-dimensional distribution functions f_1 and f_2 were calculated for laser intensities pulled from a uniform distribution between $a_0 = 4$ and $a_0 = 20$. $f_1(\langle \epsilon_f \rangle, \epsilon_{\text{crit}})$, shown in figure 3(b), agrees well with the previous work, with the classical model predicting substantially higher critical energies as expected. The predictions from the quantum and semi-classical models strongly overlap, however, and hence using critical energy from the photon spectra is a poor way of determining between stochastic and semi-classical models of radiation reaction.

Figure 3(c), however, shows another possible measurement using the same experimental parameters, comparing the mean final electron energy with the final energy spread, as described by $f_2(\langle \epsilon_f \rangle, \sigma_{\epsilon_f})$. Using these measurements the semi-classical and classical models predict fairly similar results, but the quantum model predicts a substantially higher energy spread than both of the other models. This is because, for an interaction where the electron beam is much smaller than the focal spot, these deterministic models predict that electron beams can only ever become cooled by emitting synchrotron radiation in the electric field of the laser, as more energetic electrons emit radiation more strongly. In the stochastic model, however, the number and energy of photons emitted by each electron is probabilistic and varies strongly. In certain cases, the final energy spread of the electron may increase over time (see for instance [27, 28, 34]). The lower the energy spread, the more likely this becomes. Reference [27] equation (3.8) predicts that for Gaussian energy spectra the cross-over point, below which the stochastic broadening dominates, can be approximated in the case $\sigma_{\epsilon_i} \ll \langle \epsilon_i \rangle$ by:

$$\left(\frac{\sigma_{\epsilon_i}}{\langle \epsilon_i \rangle} \right)^2 \approx \frac{55\sqrt{3}}{8 \cdot 24} \langle \chi_e \rangle, \quad (2)$$

$$\frac{\sigma_{\epsilon_i}}{\langle \epsilon_i \rangle} \approx 0.70 \sqrt{\langle \chi_e \rangle}. \quad (3)$$

We can quantify the ease of distinguishing between models by measuring the overlap between the joint distribution functions. The probability of making a measurement of several parameters, denoted by the vector \mathbf{x} , given

Table 1. Overlap Ω of joint distribution functions f_1 and f_2 and corresponding minimum number of shots required to obtain 3σ confidence in determining between models for radiation reaction, using different sets of measurements.

Models	Ω		N_{\min} for $p = 0.3\%$	
	$f_1(\langle\epsilon_f\rangle, \epsilon_{\text{crit}})$	$f_2(\langle\epsilon_f\rangle, \sigma_{\epsilon_f})$	f_1	f_2
Quantum / classical	0.235	0.134	4	3
Quantum / semi-classical	0.920	0.241	70	4
Classical / semi-classical	0.180	0.446	4	7

that a model A is true, is denoted $P(\mathbf{x}|A)$. The chances of incorrectly inferring model B from those measurements can be related by Bayes' theorem to the model probability as $P(B|\mathbf{x}) = P(\mathbf{x}|B)P(B)/P(\mathbf{x})$. If the prior assumption is that the two models are equally likely, but not the only two possible models, and that all measurements of \mathbf{x} in the region of interest are equally likely, with no bias in the measuring equipment, we can show that the probability of incorrectly inferring model B from a single measurement of \mathbf{x} , if model A is true (or vice versa) is proportional to the overlap Ω between the models, as:

$$P(B|A) = P(A|B) = \Omega \equiv \frac{\int P(\mathbf{x}|A) \cdot P(\mathbf{x}|B) dV_{\mathbf{x}}}{\sqrt{\int P(\mathbf{x}|A)^2 dV_{\mathbf{x}} \int P(\mathbf{x}|B)^2 dV_{\mathbf{x}}}}, \quad (4)$$

where the integrations are performed over the domain of possible measurements within which $P(\mathbf{x})$ is constant. This probability is normalised such that if $P(\mathbf{x}|A) = P(\mathbf{x}|B)$ then $\Omega = 1$. Depending on the choice of measurements, $P(\mathbf{x}|A)$ and $P(\mathbf{x}|B)$ are described by the joint distribution functions $f_1(\langle\epsilon_f\rangle, \epsilon_{\text{crit}})$ and $f_2(\langle\epsilon_f\rangle, \sigma_{\epsilon_f})$ for each model.

If N independent and identically distributed measurements are taken, one for each successful laser-beam collision, the probability of incorrectly inferring model B , given that A is in fact true, becomes $P(B|A) = \Omega^N$. Conversely, if we require better than a certain degree of accuracy to be sure we will not incorrectly infer model B , such that $P(B|A) < p$, we can show that we require $N > \log p / \log \Omega$. If the overlap between the joint distribution functions is very high, $\Omega \rightarrow 1$, $|\log \Omega| \rightarrow 0$ and it becomes increasingly difficult to confirm which model is correct, requiring an ever larger number of shots.

This is shown in table 1 for the parameters described in [13] and joint distribution functions shown in figure 3. The classical model predicts significantly different results to the quantum and semi-classical models and hence the overlap and number of shots required are both small. The work in [13] was therefore able to show that the quantum model agreed better with the data than the classical model, despite only definitely measuring four successful collisions. In general, measurements of ϵ_{crit} and $\langle\epsilon_f\rangle$ are successful at determining

between classical and quantum/semi-classical models of radiation reaction. With those measurements, however, it would be more difficult to distinguish between the quantum and semi-classical models, with at least 70 shots required to obtain the same level of certainty.

An alternative approach is to measure σ_{ϵ_f} , therefore significantly reducing both the overlap between the models and the number of shots required. Although an accurate measurement of σ_{ϵ_f} is difficult, requiring a clean electron energy spectrum, the difference in the predictions from quantum and semi-classical models is significant. For certain experimental parameters, using this measurement could reduce by more than an order of magnitude the number of shots required to confidently determine which model is more appropriate.

4. Optimal parameters

The overlap between the joint distribution functions assuming quantum and semi-classical models of radiation reaction was tabulated over a wide range of different initial electron energies and laser intensities in order to determine the number of shots required to distinguish at the $p = 0.3\%$ level between the two models given certain experimental parameters. Similar values as before for the uncertainties in a_0 and $\langle\epsilon_i\rangle$ were used, at $\pm 3\%$ and $\pm 10\%$ respectively, with a very large energy spread of $\sigma_{\epsilon_i}/\langle\epsilon_i\rangle = 50\%$ as before. In order to describe realistic experiments, shot-to-shot variation on σ_{ϵ_i} has also been introduced, at $\pm 25\%$ of σ_{ϵ_i} , such that the laser intensity, mean electron energy, and electron energy spread all vary shot-to-shot. The results are shown in figure 4.

Figure 4(a) shows the number of shots required when using measurements of $\langle\epsilon_f\rangle$ and ϵ_{crit} ; this demonstrates that when operating at realistic experimental parameters, many shots must be taken to distinguish between quantum and semi-classical models of radiation reaction. For low a_0 and $\langle\epsilon_i\rangle$, hundreds of shots are required to conclude that one of the models is correct and not the other. In contrast to the situation for mono-energetic electron beams, increasing the laser intensity above $a_0 = 10$ increases the difference between the predictions from the two models. If it is possible to increase the electron energy to $\langle\epsilon_i\rangle \geq 1$ GeV and the laser intensity to $a_0 \geq 15$, the number of shots required is reduced to below 25. Under these conditions, a practical radiation reaction experiment could determine with significant ($>3\sigma$) confidence which model is more appropriate in this regime.

If we cannot determine between quantum and semi-classical models of radiation reaction using measurements of $\langle\epsilon_f\rangle$ and ϵ_{crit} , it is possible to use a measurement of the final electron energy spread σ_{ϵ_f} . As shown in figure 4(b), this does not reduce the number of shots required at low laser intensities, but is substantially more successful at lower electron energies, with fewer than 25 shots required if $\langle\epsilon_i\rangle \geq 500$ MeV and $a_0 \geq 15$. For sufficiently high electron energies and laser intensities (for instance $\langle\epsilon_i\rangle \geq 1$ GeV and $a_0 \geq 20$, the predictions from the quantum and semi-classical models for radiation reaction are substantially different; under these conditions, it is vital to understand which model of radiation




Figure 4. The estimated number of shots required to distinguish between the quantum and semi-classical models at the $p = 0.3\%$ significance level using measurement of (a) $\langle \epsilon_f \rangle$ and ϵ_{crit} and (b) $\langle \epsilon_f \rangle$ and σ_{ϵ_f} , plotted against the laser intensity a_0 and the electron energy ϵ_i . The variation on a_0 was taken to be ± 3 , and the shot-to-shot variation on $\langle \epsilon_i \rangle$ was $\pm 10\%$. The initial energy spread was $\sigma_{\epsilon_i} = (0.5 \pm 0.125) \langle \epsilon_i \rangle$. The colour scale is the same for both plots.




Figure 5. Using an initial energy spread of (a) $\sigma_{\epsilon_i}/\langle \epsilon_i \rangle = 20\% \pm 5\%$ and (b) $\sigma_{\epsilon_i}/\langle \epsilon_i \rangle = 10\% \pm 2.5\%$, the estimated number of shots required to distinguish between quantum and semi-classical models of radiation reaction at a confidence level of $p = 0.3\%$. All values are using measurements of $\langle \epsilon_f \rangle$ and σ_{ϵ_f} and are plotted against the laser intensity a_0 and the electron energy ϵ_i . The shot-to-shot variation on a_0 was taken to be ± 3 , and the variation of ϵ_i was $\pm 10\%$. The colour scale is the same for both plots.

reaction is more accurate, and only a single shot may be sufficient to discriminate between quantum and semi-classical models.

It is clear that the energy spread of the electron spectrum is the key distinguishing feature of a stochastic model as opposed to a deterministic model, and we can study the effect of decreasing the initial energy spread in the electron spectrum. This will reduce the cooling experienced by the electron beam, increasing the relative contribution of stochasticity. Simulated experiments were run as before, but with the relative energy spread reduced to 20% and 10%. Again, the relative variation of the energy spread was $\pm 25\%$, giving energy spreads of $\sigma_{\epsilon_i} = (0.2 \pm 0.05) \langle \epsilon_i \rangle$ and $(0.1 \pm 0.025) \langle \epsilon_i \rangle$, respectively. As shown in figure 5, reducing the initial energy spread significantly changes the final energy spread and reduces the number of shots required to determine which model is more suitable. If no more than one hundred shots are possible and the relative energy spread is 20%, distinguishing between the models requires just $\langle \epsilon_i \rangle \geq 500$ MeV or $a_0 \geq 12$. Under most conditions simulated, fewer than 10 shots would be required. For an energy

spread of 10%, however, the models can easily be distinguished even for the lowest laser intensities and electron energies. Only at $\langle \epsilon_i \rangle = 200$ MeV and $a_0 = 5$ do the models predict very similar outcomes; under these conditions the accuracy of the constant-cross-field approximation is doubtful and it is likely that both models will break down. In most of the simulated experiments, however, only a single shot would be sufficient to determine which model is more correct. Reducing the energy spread of the initial electron beam is therefore one of the best ways of ensuring an experiment will be able to distinguish between deterministic and stochastic models of radiation reaction.

We can study the maximum allowable energy spread at a certain laser intensity and initial electron energy, if stochastic effects are to be measured. First, we ran simulated experiments for $a_0 = 10 \pm 3$ and $\langle \epsilon_i \rangle = (500 \pm 50)$ MeV, which corresponds to a quantum parameter $\langle \chi_e \rangle = \langle \gamma_e E_L / E_s \rangle \approx 0.03$. These parameters are achievable in many existing PW scale laser facilities, and are on the same scale as achieved previously in [13]. The results, plotted in figure 6(a), show the number of shots required for a range of energy spreads σ_{ϵ_i} and errors on the




Figure 6. Using initial conditions of (a) $a_0 = 10 \pm 3$ and $\langle \epsilon_i \rangle = (500 \pm 50) \text{ MeV}$, and (b) $a_0 = 15 \pm 3$ and $\langle \epsilon_i \rangle = (1.0 \pm 0.1) \text{ GeV}$, the number of shots required for Monte-Carlo simulated experiments to determine with 3σ confidence between the quantum and semi-classical models. All values are using measurements of $\langle \epsilon_f \rangle$ and σ_{ϵ_f} and are plotted against the electron energy spread, relative to the mean electron energy, and the shot-to-shot variation of the electron energy spread, relative to the energy spread. The colour scale is the same for both plots. The analytic prediction to measure stochastic broadening in the ideal case is also shown (white dotted line).

energy spread. The number of shots required increases with both the energy spread and the variation in the energy spread, and is sufficient to make experiments impractical when either of these are high. For an experiment to be practical, requiring only 10 s of shots to successfully distinguish between models, the energy spread must generally be kept below around $\sigma_{\epsilon_i}/\langle \epsilon_i \rangle \lesssim 25\%$. If the relative error on the energy spread can be greatly reduced, however, experiments with these parameters can be successful while $\sigma_{\epsilon_i}/\langle \epsilon_i \rangle \leq 50\%$. For $\langle \chi_e \rangle \approx 0.03$ in equation (3), the energy spread required for the electron spectrum to broaden is $\sigma_{\epsilon_i}/\langle \epsilon_i \rangle \lesssim 12\%$, below which simulated experiments measure a clear difference between quantum and semi-classical models. However, the simulated experiments show that a significant difference arises between the two models well before stochastic broadening dominates, so long as the variation on the energy spread is limited to a few tens of percent or lower.

We also ran simulated experiments for $a_0 = 15 \pm 3$ and $\langle \epsilon_i \rangle = (1.0 \pm 0.1) \text{ GeV}$, or $\langle \chi_e \rangle \approx 0.1$, parameters which are around the limit of what is achievable with some of the existing petawatt laser facilities (e.g. [14]). Figure 6(b) shows the results, demonstrating that the number of shots required is greatly reduced at these parameters. Under these conditions, it is relatively straightforward to distinguish between the quantum and semi-classical models, with the predicted energy spread significantly different even when the energy spread is on the level of $\sigma_{\epsilon_i}/\langle \epsilon_i \rangle \approx 50\%$ and varies widely shot-to-shot. Again, we can calculate where stochastic broadening dominates analytically using equation (3), giving a condition on the energy spread of $\sigma_{\epsilon_i}/\langle \epsilon_i \rangle \lesssim 20\%$. With this increased laser intensity and electron energy, the simulations show that the predictions of the two models diverge even well above this threshold, with the stochastic model predicting a large reduction in the cooling rate and a significantly different final energy spread.

Finally, we explored the effect of shot-to-shot variation of the laser intensity and mean initial electron energy under the same two sets of conditions: $a_0 = 10$ and

$\langle \epsilon_i \rangle = 500 \text{ MeV}$; and $a_0 = 15$ and $\langle \epsilon_i \rangle = 1 \text{ GeV}$. The energy spread was set as $\sigma_{\epsilon_i} = 0.25\langle \epsilon_i \rangle$, with the shot-to-shot variation on the energy spread $0.25\sigma_{\epsilon_i}$. Figure 7 shows the number of shots required at a range of errors on both a_0 and $\langle \epsilon_i \rangle$. Interestingly, the effect is small, with no drastic change in the overlap between predictions from the two different models. At a higher laser intensity and electron energy the number of shots required increases slowly with an increasing shot-to-shot variation, as expected. At the lower intensity and electron energy, however, large variation actually results in more significant radiation reaction effects in the high energy and high intensity tails, causing a slight reduction in the number of shots required. Experiments will remain able to distinguish between stochastic and deterministic models of radiation reaction even if the shot-to-shot variation in the laser intensity and electron energy are high, so long as the shot-to-shot variation in electron energy spread is limited.

In the course of these simulated experiments we have considered a wide range of experimental errors, but the laser pulse profile and the electron energy distribution have remained idealised. In practice, electron bunches from LWFA often contain a significant lower energy or thermal component, particularly when operating in the so-called ‘bubble’ regime [35], where electrons are continuously injected into the wakefield. This work has neglected that background, which would have to be carefully removed from the energy spectrum before measuring the final energy spread of the beam. Spatial variation in both the electron beam and the laser pulse can also result in significant changes to the final electron energy spread; if one region of the electron beam experiences a much higher laser intensity and a greater radiation reaction force, the final energy spread of the beam can be significantly higher than expected. Practical experiments must work to limit the spatial variation, which can be achieved by moving the focal plane of the laser pulse further away from the point of collision with the electron beam, at the cost of reducing the effective laser intensity. Future simulated experiments, on the




Figure 7. For (a) $a_0 = 10$ and $\langle \epsilon_i \rangle = 500 \text{ MeV}$, and (b) $a_0 = 15$ and $\langle \epsilon_i \rangle = 1 \text{ GeV}$, the estimated number of shots required to distinguish at the $p = 0.3\%$ significance level between the quantum and semi-classical models. All values are using measurements of $\langle \epsilon_f \rangle$ and σ_{ϵ_f} and are plotted against the relative shot-to-shot variation on the initial electron energy and the variation on the laser intensity. The colour scale is the same for both plots.

other hand, could include variation of the laser intensity within a single shot, as well as shot-to-shot [36].

5. Conclusions

We have used a series of particle-in-cell simulations and Monte-Carlo simulated experiments to make predictions for radiation reaction from each of the quantum, classical, and semi-classical models using realistic parameters. In doing so, we have shown that while measurements correlating the critical energy of the resulting photon spectra with the mean final electron energy give a way of clearly distinguishing the classical model from the quantum and semi-classical models, this is a poor way of determining which of the quantum and semi-classical models to use. For laser intensities $a_0 < 15$ and electron energies $\epsilon_i < 1 \text{ GeV}$, these two models predict almost the same final average energy of electrons and photons, for both mono-energetic electron beams and more realistic broad electron distributions.

Instead, we have shown that measuring the energy spread of the electron spectrum after the interaction gives a clearer distinction between the stochastic and deterministic models. Although the energy spread will only increase over the course of the interaction when the initial energy spread is very low (around $\sigma_{\epsilon_i}/\langle \epsilon_i \rangle \leq 12\%$ for $a_0 = 10$ and $\epsilon_i = 500 \text{ MeV}$), the effect of stochasticity substantially reduces the rate of cooling, leading to different predictions from the quantum and semi-classical models even at much higher energy spreads.

We have used the simulated experiments to determine how many shots would be required when operating at certain conditions, and used this to build up a picture of the optimal experimental parameters. Crucially electron energy spread should be reduced to below 25%, if possible, to maximise the chances of conclusively determining which model is more accurate, while shot-to-shot variation on the energy spread should be minimised. At this energy spread, it should be possible to distinguish the quantum and semi-classical models in a few shots even with relatively unambitious experimental

parameters, such as $\langle \epsilon_i \rangle = 500 \text{ MeV}$ and $a_0 = 10$. Under these conditions, the measurement is robust to significant shot-to-shot variation in electron energy and laser intensity, even on the scale of 50%.

Alternatively, if it is difficult to reduce the electron energy spread, increasing the electron energy or laser intensity will separate the predictions from the two models. Even with an energy spread of $\sigma_{\epsilon_i} = 50\%$, it is possible to distinguish between the models fairly clearly if $a_0 > 10$ and $\langle \epsilon_i \rangle > 1 \text{ GeV}$, or if $a_0 > 15$ and $\langle \epsilon_i \rangle > 500 \text{ MeV}$. By an electron energy of $\langle \epsilon_i \rangle = 1 \text{ GeV}$ and a laser intensity of $a_0 = 15$, the models can be fairly easily distinguished regardless of a large energy spread or shot-to-shot variation. These requirements are certainly achievable using current laser systems, and upcoming experiments should be able to clearly determine which model of radiation reaction is most suitable for describing interactions of energetic electrons with high intensity lasers.

Acknowledgments

CA and CPR are grateful for EPSRC Grant No. EP/M018156/1 which made this work possible. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 682399), from EPSRC Grant No. EP/M018555/1, and from the Knut and Alice Wallenberg Foundation. Simulations used the EPOCH PIC code developed under UK EPSRC Grants EP/G054950/1, EP/G056803/1, EP/G055165/1 and EP/M022463/1, while computing resources were provided by STFC Scientific Computing Department's SCARF cluster.

Appendix A. EPOCH revisions

The quantum electro-dynamics model in EPOCH described in detail in [22] is a Monte-Carlo stochastic model which takes

into account both the changes to the synchrotron spectrum and emission rate as the effective electric field rises, and also the random nature of the photon emissions. A charged macro-particle is initialised with a random optical depth τ_0 . Its optical depth is then reduced at every timestep by:

$$\tau_n = \tau_{n-1} - \frac{dN_\gamma}{dt} \cdot \delta t, \quad (\text{A1})$$

where τ_n is the optical depth at the n th timestep and δt is the duration of the timestep. When the optical depth falls below zero, the macro-particle emits a single macro-photon, with a particle weight equal to that of the original macro-particle and an energy chosen at random from the relevant synchrotron spectrum, as:

$$\varepsilon_\gamma = 2m_e c^2 \gamma_e \frac{\chi_\gamma}{\chi_e}, \quad (\text{A2})$$

$$w_\gamma = w_e, \quad (\text{A3})$$

where χ_γ is chosen from the synchrotron spectrum $P_{\chi_\gamma}(\chi_\gamma|\chi_e)$, where $\chi_e = \gamma_e E_L/E_s$, where $\gamma_e = 1 + \varepsilon_e/m_e c^2$ is the usual relativistic factor, and $E_s \approx 1.32 \times 10^{18} \text{ V m}^{-1}$ is the Schwinger limit.

We wish, however, to explore two alternative deterministic models: the fully classical model, which possesses neither the random nature of emission nor the changes to the synchrotron spectrum and emission rate; and the so-called semi-classical model, which contains the changes to the synchrotron spectrum and rate, but not the random emission. In these two models, each macro-electron now emits a macro-photon at every timestep, ignoring the optical depth. The energy is chosen at random from the relevant synchrotron spectrum as before (with the classical model using the limit of the synchrotron spectrum as $\chi_e \rightarrow 0$), but the particle weight of this macro-photon is now proportional to the instantaneous emission rate and the timestep duration as:

$$w_\gamma = w_e \cdot \frac{dN_\gamma}{dt} \cdot \delta t \quad (\text{A4})$$

In this way, the semi-classical model will predict the same photon energy spectrum and rate of emission as the quantum model, but the semi-classical model is deterministic, with no element of randomness. In this model, charged particles continually emit photons.

For vanishing small $E_L \rightarrow 0$ both the particle weight and energy of the macro-photon should vanish to zero and emission under these conditions will contribute negligibly to the final photon spectrum. There is, however, an additional complexity due to the implementation of $P_{\chi_\gamma}(\chi_\gamma|\chi_e)$ in EPOCH. This is tabulated, and has a lower limit as χ_e vanishes to zero with $E_L \rightarrow 0$, at a value $\chi_{e,\min}$. This implies that as χ_e vanishes to zero, χ_γ does not, and so $\varepsilon_\gamma \rightarrow \infty$, which is clearly unphysical.

As $\chi_e \rightarrow 0$, $P_{\chi_\gamma}(\chi_\gamma|\chi_e)$ should instead tend towards the classical limit, where the synchrotron spectrum is a function of a single variable only: $P_{\chi_\gamma}(\chi_\gamma|\chi_e) \rightarrow P(\chi_\gamma/\chi_e^2)$. For

$\chi_e < \chi_{e,\min}$, we instead calculate the photon energy using:

$$\chi_\gamma = \chi'_\gamma \left(\frac{\chi_e}{\chi_{e,\min}} \right)^2, \quad (\text{A5})$$

where χ'_γ is chosen at random from $P_{\chi_\gamma}(\chi'_\gamma|\chi_{e,\min})$. In this way, even though χ'_γ cannot vanish to zero, both χ_γ and $\varepsilon_\gamma \propto \chi_\gamma/\chi_e$ will safely vanish to zero as $\chi_e \rightarrow 0$.

This step is not generally necessary in the stochastic case, as the probability of emitting a macro-photon vanishes to zero as $E_L \rightarrow 0$, so photons with un-physically high energies are never created. When emitting a macro-photon at every timestep, however, this step is important to avoid a large population of extremely high energy macro-photons, even though the particle weights of these macro-photons safely vanish to zero.

Appendix B. Monte-Carlo simulated experiments

For each simulated shot, values for the laser intensity, parameterised by $a_0^{(n)}$, the mean initial electron energy $\langle \varepsilon_i \rangle^{(n)}$, and the initial energy spread $\sigma_{\varepsilon_i}^{(n)}$ were randomly chosen from Gaussian distributions with a chosen mean and standard deviation. $N_{\text{electron}} = 10\,000$ electrons were then simulated, each with an initial energy $\varepsilon_i^{(s)}$ sampled from another Gaussian distribution, with mean $\langle \varepsilon_i \rangle^{(n)}$ and standard deviation $\sigma_{\varepsilon_i}^{(n)}$. For each simulated electron, the final energy distribution was characterised by a Gaussian with mean $\langle \varepsilon_f \rangle(\varepsilon_i^{(s)}, a_0^{(n)})$ and standard deviation $\sigma_{\varepsilon_f}(\varepsilon_i^{(s)}, a_0^{(n)})$ drawn from the look-up table, and the final electron energy was then estimated by drawing a random sample $\varepsilon_f^{(s)}$ from this distribution. Example electron energy spectra from the mono-energetic simulations are shown in figure B1(a) alongside the Gaussian fits. When the electron energy and laser intensity are very high, the final electron spectra are strongly skewed and the Gaussian distribution becomes a worse approximation, but at the laser intensities and electron energies considered in this paper the divergence is small.

A histogram was assembled for each shot, using the N_{electron} different values of $\varepsilon_f^{(s)}$, and this was fitted to a Gaussian to estimate $\langle \varepsilon_f \rangle^{(n)}$ and $\sigma_{\varepsilon_f}^{(n)}$. At the same time, for each electron the photon distribution $N_\gamma(\varepsilon_\gamma|\varepsilon_i^{(s)}, a_0^{(n)})$ was calculated from the look-up table and added to a total distribution $N_\gamma^{(n)}(\varepsilon_\gamma|N_{e,i}^{(n)}, a_0^{(n)})$. This was then fitted to equation (1) to give an estimate $\varepsilon_{\text{crit}}^{(n)}$. Examples of photon spectra from the mono-energetic simulations are shown in figure B1(b) alongside fits to equation (1). Summing the photon spectra means that, as in real experiments, the much higher number of photons emitted by the highest energy electrons tend to dominate the spectrum and using $\varepsilon_{\text{crit}}^{(n)}$ remains a reasonable way of parameterising the measured spectrum.

In total, $N_{\text{shots}} = 10\,000$ shots were simulated, and the different estimates of $\langle \varepsilon_f \rangle^{(n)}$, $\sigma_{\varepsilon_f}^{(n)}$ and $\varepsilon_{\text{crit}}^{(n)}$ were combined using a Gaussian kernel-density estimate to form the joint distribution function $f(\varepsilon_{\text{crit}}, \langle \varepsilon_f \rangle, \sigma_{\varepsilon_f})$. The process was repeated for each of the classical, semi-classical and quantum models of radiation reaction, using the appropriate look-up tables, giving three different joint distribution functions $f^{(Q)}$,

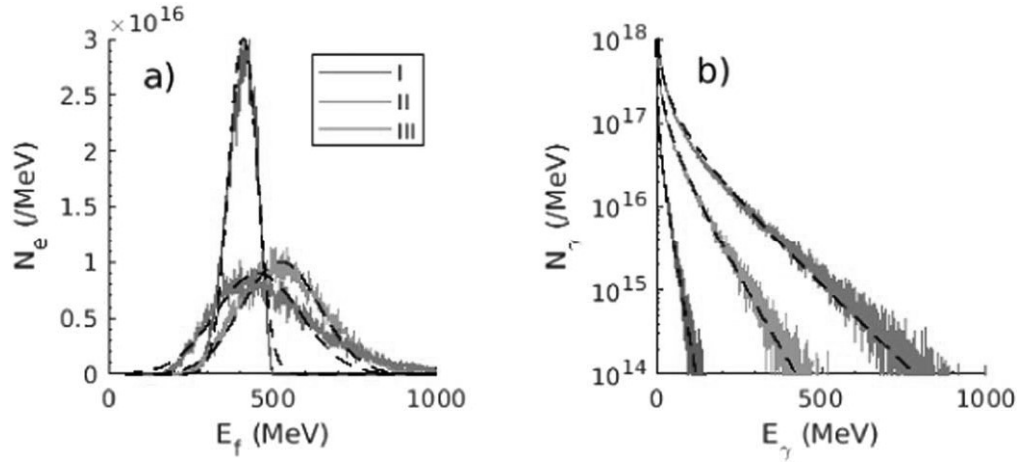


Figure B1. Example spectra for (a) electrons and (b) photons resulting from simulations of mono-energetic electron beams interacting with laser pulses. Electron spectra were fitted to Gaussian distributions (dashed black lines), whereas photon spectra were fitted to equation (1). Conditions were (I) a low laser intensity and electron energy ($a_0 = 10$, $E_i = 500$ MeV), (II) a moderate intensity and energy ($a_0 = 15$, $E_i = 1000$ MeV), and (III) a high intensity and energy ($a_0 = 20$, $E_i = 1500$ MeV).

$f^{(S)}$, and $f^{(C)}$. For parameter scans, however, it is computationally expensive to calculate the three-dimensional distribution function, and so the two-dimensional distribution functions $f_1(\langle \varepsilon_f \rangle, \varepsilon_{\text{crit}})$ and $f_2(\langle \varepsilon_f \rangle, \sigma_{\varepsilon_f})$ were used. These are effectively integrations of $f(\varepsilon_{\text{crit}}, \langle \varepsilon_f \rangle, \sigma_{\varepsilon_f})$, integrated over σ_{ε_f} or $\varepsilon_{\text{crit}}$ respectively, and require just $N_{\text{shots}} = 1000$ shots to accurately sample the underlying distribution. Convergence testing, varying N_{shots} , allowed us to estimate the error on the probability under these conditions as approximately 5%.

Appendix C. Optimal pulse duration

In order to study the optimal pulse duration we can consider the stochastic contribution to broadening of the electron energy distribution. This is described by the second moment of the synchrotron emission distribution function (or the first moment of the synchrotron energy spectrum) [27]:

$$g_2(\chi_e) \equiv \frac{\int_0^{\chi_e/2} \chi_\gamma F(\chi_e, \chi_\gamma) d\chi_\gamma}{\int_0^\infty \chi_\gamma F_{\text{cl}}(4\chi_\gamma/3\chi_e^2) d\chi_\gamma}, \quad (\text{C1})$$

where $F(\chi_e, \chi_\gamma)$ and $F_{\text{cl}}(4\chi_\gamma/3\chi_e^2)$ are the quantum and classical synchrotron functions respectively, which describe the energy spectra [21]. The function $g_2(\chi_e)$ can be approximated by $g_2(\chi_e) \approx [1 + (1 + 4.528\chi_e) \ln(1 + 12.29\chi_e) + 4.632\chi_e^2]^{-7/6}$.

Stochastic broadening results in a rate of increase in the variance $\sigma^2 \equiv \langle \gamma^2 \rangle - \langle \gamma \rangle^2$, described by:

$$\left(\frac{d\sigma^2}{dt} \right)_+ = \frac{\langle S \rangle}{m_e^2 c^4}, \quad (\text{C2})$$

where

$$S(\chi_e) = \frac{55\alpha}{24\sqrt{3}\tau_c} m_e^2 c^4 \gamma \chi_e^3 g_2(\chi_e), \quad (\text{C3})$$

and the average $\langle S \rangle$ is taken over the electron population.

For a laser pulse with a duration τ and a total energy E , $\chi_e \propto \gamma a_0 \propto \gamma E^{\frac{1}{2}} \tau^{-\frac{1}{2}}$ and the total increase in the energy spread due to this stochastic emission is therefore described by:

$$\Delta\sigma^2 \propto \gamma \chi_e^3 g_2(\chi_e) \tau \propto E \gamma^3 \chi_e g_2(\chi_e). \quad (\text{C4})$$

At low $\chi_e \ll 1$, $g_2(\chi_e) \approx 1$, and for a fixed laser energy, broadening only increases as the pulse duration falls, as $\Delta\sigma^2 \propto E^{\frac{3}{2}} \gamma^4 \tau^{-\frac{1}{2}}$. However, for $\chi_e \gg 1$, $g_2(\chi_e)$ falls towards zero as $g_2(\chi_e) \propto \chi_e^{-\frac{7}{3}}$ and the increase in energy spread also falls towards zero as $\Delta\sigma^2 \propto E \gamma^3 \chi_e^{-\frac{4}{3}} \propto E^{\frac{1}{3}} \gamma^{\frac{5}{3}} \tau^{\frac{2}{3}}$.

The increase in energy spread is maximal in between these two regimes, where:

$$\frac{\partial \Delta\sigma^2}{\partial \chi_e} = 0 = g_2(\chi_e) + \chi_e \frac{\partial g_2(\chi_e)}{\partial \chi_e}. \quad (\text{C5})$$

Solving this numerically gives $\chi_e \approx 0.56$, which for a ≈ 2 GeV electron bunch corresponds to $a_0 \approx 50$. This is beyond the regime explored by the radiation reaction experiments considered in this paper, but well within the capabilities of 10 PW laser facilities. The results of this paper, considering only $a_0 \leq 25$, suggest increasing the laser intensity in order to better distinguish between models of radiation reaction, but experiments using 10 PW laser facilities could consider using laser pulses which are not fully compressed, with lower intensities than the maximum possible laser intensity, in order to increase the broadening of the electron energy distribution in future radiation reaction experiments.

ORCID iDs

C Arran <https://orcid.org/0000-0002-8644-8118>

E Gerstmayr <https://orcid.org/0000-0003-1164-8593>

T G Blackburn <https://orcid.org/0000-0002-3681-356X>

References

- [1] Zou J *et al* 2015 *High Power Laser Sci. Eng.* **3** 2
- [2] Weber S *et al* 2017 *Matter Radiat. Extremes* **2** 149

- [3] Gales S *et al* 2018 *Rep. Prog. Phys.* **81** 094301
- [4] Bell A R and Kirk J G 2008 *Phys. Rev. Lett.* **101** 200403
- [5] Di Piazza A, Müller C, Hatsagortsyan K Z and Keitel C H 2012 *Rev. Mod. Phys.* **84** 1177
- [6] Mourou G A, Tajima T and Bulanov S V 2006 *Rev. Mod. Phys.* **78** 309
- [7] Powers N D, Ghebregziabher I, Golovin G, Liu C, Chen S, Banerjee S, Zhang J and Umstadter D P 2013 *Nat. Photon.* **8** 28
- [8] Tsai H E *et al* 2016 *Phys. Plasmas* **22** 023106
- [9] Duff M J, Capdessus R, Sorbo D D, Ridgers C P, King M and McKenna P 2018 *Plasma Phys. Control. Fusion* **60** 064006
- [10] Tamburini M, Pegoraro F, Piazza A D, Keitel C H and Macchi A 2010 *New J. Phys.* **12** 123005
- [11] Zhang P, Ridgers C P and Thomas A G R 2015 *New J. Phys.* **17** 043051
- [12] Wang W-M, Gibbon P, Sheng Z-M, Li Y-T and Zhang J 2017 *Phys. Rev. E* **96** 013201
- [13] Cole J M *et al* 2018 *Phys. Rev. X* **8** 011020
- [14] Poder K *et al* 2018 *Phys. Rev. X* **8** 031004
- [15] Tajima T and Dawson J 1979 *Phys. Rev. Lett.* **43** 267
- [16] Mangles S P D, Murphy C D and Najmudin Z 2004 *Nature* **431** 535
- [17] Faure J, Glinec Y, Pukhov A, Kiselev S, Gordienko S, Lefebvre E, Rousseau J-P, Burgy F and Malka V 2004 *Nature* **431** 541
- [18] Geddes C G R, Toth C S, Van Tilborg J, Esarey E, Schroeder C B, Bruhwiler D, Nieter C, Cary J and Leemans W P 2004 *Nature* **431** 538
- [19] Schwinger J 1951 *Phys. Rev.* **82** 664
- [20] Erber T 1966 *Rev. Mod. Phys.* **38** 626
- [21] Sokolov I V, Nees J A, Yanovsky V P, Naumova N M and Mourou G A 2010 *Phys. Rev. E* **81** 036412
- [22] Ridgers C, Kirk J, Duclous R, Blackburn T, Brady C, Bennett K, Arber T and Bell A 2014 *J. Comput. Phys.* **260** 273
- [23] Blackburn T G, Ridgers C P, Kirk J G and Bell A R 2014 *Phys. Rev. Lett.* **112** 015001
- [24] Landau L D and Lifshitz E M 1971 *The Classical Theory of Fields*, vol 2 3rd edn (Oxford: Pergamon)
- [25] Macchi A 2018 *APS Phys.* **11** 13
- [26] Neitz N and Piazza A Di 2013 *Phys. Rev. Lett.* **111** 054802
- [27] Ridgers C P *et al* 2017 *J. Plasma Phys.* **83** 715830502
- [28] Niel F, Riconda C, Amiranoff F, Duclous R and Grech M 2018 *Phys. Rev. E* **97** 043209
- [29] Ritus V I 1985 *J. Russ. Laser Res.* **6** 497
- [30] Harvey C N, Ilderton A and King B 2015 *Phys. Rev. A* **91** 013822
- [31] Di Piazza A, Tamburini M, Meuren S and Keitel C H 2018 *Phys. Rev. A* **98** 012134
- [32] Blackburn T G, Seipt D, Bulanov S S and Marklund M 2018 *Phys. Plasmas* **25** 083108
- [33] Arber T D *et al* 2015 *Plasma Phys. Control. Fusion* **57** 113001
- [34] Vranic M, Grismayer T, Fonseca R A and Silva L O 2016 *New J. Phys.* **18** 073035
- [35] Pukhov A M, Gordienko S, Kiselev S and Kostyukov I 2004 *Plasma Phys. Control. Fusion* **46** B179
- [36] Baird C D, Murphy C D, Blackburn T G, Ilderton A, Mangles S P D, Marklund M and Ridgers C P 2019 *New J. Phys.* **21** 053030