Laser wakefield acceleration with active feedback at 5 Hz

S. J. D. Dann,[1,2,*] C. D. Baird,[3] N. Bourgeois,[2] O. Chekhlov,[2] S. Eardley,[4] C. D. Gregory,[2] J.-N. Gruse,[5] J. Hah,[6] D. Hazra,[7] S. J. Hawkes,[2] C. J. Hooker,[2] K. Krushelnick,[6] S. P. D. Mangles,[5] V. A. Marshall,[2] C. D. Murphy,[3] Z. Najmudin,[5] J. A. Nees,[6] J. Osterhoff,[8] B. Parry,[2] P. Pourmoussavi,[8] S. V. Rahul,[9] P. P. Rajeev,[2] S. Rozario,[5] J. D. E. Scott,[1] R. A. Smith,[4] E. Springate,[2] Y. Tang,[2] S. Tata,[10] A. G. R. Thomas,[1,6] C. Thornton,[2] D. R. Symes,[2] and M. J. V. Streeter[1,5]

[1]*Physics Department, Lancaster University, Bailrigg, Lancaster LA1 4YW, United Kingdom*
[2]*Central Laser Facility, Rutherford Appleton Laboratory, Didcot OX11 0QX, United Kingdom*
[3]*York Plasma Institute, University of York, Heslington, York YO10 5DQ, United Kingdom*
[4]*Blackett Laboratory, Imperial College London, London SW7 2AZ, United Kingdom*
[5]*The John Adams Institute for Accelerator Science, Imperial College London, London SW7 2AZ, United Kingdom*
[6]*Center for Ultrafast Optical Science, University of Michigan, Ann Arbor, Michigan 48109-2099, USA*
[7]*Laser Plasma Section, Raja Ramanna Centre for Advanced Technology, Indore 452013, India*
[8]*Deutsches Elektronen-Synchrotron DESY, Notkestrasse 85, 22607 Hamburg, Germany*
[9]*TIFR Centre for Interdisciplinary Sciences, Hyderabad 500 107, India*
[10]*Tata Institute of Fundamental Research, Homi Bhabha Road, Colaba, India*

(Received 20 December 2018; published 26 April 2019)

We describe the use of a genetic algorithm to apply active feedback to a laser wakefield accelerator at a higher power (10 TW) and a lower repetition rate (5 Hz) than previous work. The temporal shape of the drive laser pulse was adjusted automatically to optimize the properties of the electron beam. By changing the software configuration, different properties could be improved. This included the total accelerated charge per bunch, which was doubled, and the average electron energy, which was increased from 22 to 27 MeV. Using experimental measurements directly to provide feedback allows the system to work even when the underlying acceleration mechanisms are not fully understood, and, in fact, studying the optimized pulse shape might reveal new insights into the physical processes responsible. Our work suggests that this technique, which has already been applied with low-power lasers, can be extended to work with petawatt-class laser systems.

DOI: 10.1103/PhysRevAccelBeams.22.041303

I. INTRODUCTION

Plasma waves are capable of supporting electromagnetic fields which are orders of magnitude stronger than can be maintained in radio-frequency cavities. In a laser-driven plasma wakefield accelerator (LWFA), an intense laser pulse is used to drive a plasma wave, which then propagates at close to the speed of light. By using a driving laser pulse with a short duration compared to the typical plasma wavelength $\tau_L < 2\pi/\omega_p$, where ω_p is the plasma frequency, the plasma wave is strongly driven [1]. With a relativistically intense laser pulse, i.e., normalized vector potential $a_0 \gtrsim 1$, the peak electrostatic fields in the plasma wave are of the order of the critical electric field strength for wave breaking in a cold plasma [2], $E_{\max} \sim E_{\mathrm{crit}} = \omega_p m_e c/e$. For typical electron densities $n_e \sim 10^{18}$ cm^{-3}, $\omega_p = 5.6 \times 10^{13}$ fs^{-1}, and so the peak accelerating field is on the order of 100 GeV m^{-1}. Using modern laser systems, LWFAs are capable of producing narrow energy spread electron beams [3–5] and maximum beam energies in excess of 1 GeV [6,7] in just a few centimeters.

One of the most promising applications for the LWFA is as a compact x-ray source. LWFA-generated electron beams have been used to generate radiation in magnetic undulators [8,9], and several projects are underway worldwide to drive a free-electron laser, using plasma-accelerated beams. In addition, the strong focusing forces in the plasma wakefield structure cause the electron beam to oscillate transversely during acceleration with a typical betatron frequency $\omega_\beta = \omega_p/\sqrt{2\gamma}$, where $\gamma m_e c^2$ is the electron energy. In this way, a bright source of x rays is produced

*stephen.dann@stfc.ac.uk

with a small source size and temporal duration determined by the radiating electron bunch [10–12]. This plasma wiggler source has been successfully used for imaging purposes [13,14] and offers the possibility to create miniature synchrotronlike facilities [15].

The average power of such sources is currently limited by the low repetition rates of commonly available petawatt lasers. However, new technologies such as diode-pumped solid state lasers can already operate at 10 Hz [16] and will soon be used to pump petawatt lasers [17]. In addition to increased average power, these higher repetition rates will bring other opportunities. One of these is the possibility of using active feedback to directly optimize the properties of the electron beam. Better control over the quality and consistency of the electron beams is needed to develop practical, compact radiation sources.

Using feedback to optimize laser-matter interactions is a well-established technique, for example, in the coherent control of atomic and molecular processes [18,19]. Applying the same principle to laser wakefield acceleration is a more recent development. He *et al.* successfully optimized various beam parameters, including the charge, divergence, and energy, by adjusting the laser wavefront using a deformable mirror [20,21]. The laser in that case produced pulse energies of 15 mJ at 500 Hz. In this work, we apply similar methods to a laser system with a significantly higher pulse energy, thus accessing an LWFA regime at the tens of MeV level. By demonstrating the concept at a lower repetition rate, we prove that this type of active feedback technique will be applicable to the 10 Hz, petawatt laser systems becoming available in the near future.

Various different approaches can be taken to experimental optimization through active feedback. One approach, which works even for cases where the underlying physics is not fully understood, is to treat it as a numerical optimization problem. We have a function of n variables, called the goal function, and aim to find the point in the n-dimensional parameter space at which the value of the goal function is maximized. Algorithms to solve this problem work by evaluating the goal function, and sometimes its derivative, at a number of points within the parameter space, eventually converging to an optimum. Following the lead of He *et al.*, we have concentrated on using genetic algorithms, but many others exist.

In this case, the goal function includes the entire experiment. Its inputs represent certain experimental settings that control the spatial and temporal structure of the driving laser pulse. The process of evaluating the goal function involves configuring the experiment according to the inputs, collecting data, and then analyzing that data to give a single figure of merit. It is that figure of merit that serves as the output of the function, which will be maximized.

In most applications of numerical optimization, the goal function is a mathematical formula or computer program, and its values are calculated. Here, the goal function is a physical process, and its values must be measured. This leads to certain challenges: Measurements are affected both by experimental errors and by longer-term drifts in experimental conditions. The measurement process also takes a relatively long time, especially as multiple signals must be averaged to reduce measurement errors. And, unlike a long-running calculation, it cannot be run on multiple machines in parallel or left running unattended. These problems have certain implications for the design of the optimization algorithm, which must be able to deal with noisy data and converge after a relatively small number of function evaluations.

In this paper, we describe an experiment in which we applied active feedback to optimize the properties of the electron beam produced by a laser wakefield accelerator. The laser used had a peak power of 15 TW and a repetition rate of 5 Hz. To the best of our knowledge, this is the first time active feedback has been applied to directly control laser wakefield acceleration at such a laser power. In Sec. II, we describe the experimental setup. In Sec. III, we describe the optimization algorithms used for active feedback. In Sec. IV, we present the results of selected optimization runs, demonstrating significant improvements in the accelerated charge and some degree of control over the beam energy. These are discussed in more detail in Sec. V. We compare the performance of two different optimization algorithms in Sec. VI. Finally, in Sec. VII, we summarize our main conclusions and present ideas for future work.

II. EXPERIMENTAL METHODS

The experiment, shown schematically in Fig. 1, was performed using the 5 Hz section (target area 2) of the Gemini laser system at the Central Laser Facility. The laser pulses, with 450 mJ on-target energy and 40 fs FWHM duration, were focused onto the edge of a 3-mm-diameter methane gas jet (Parker series 9) by an $f/16$ off-axis parabolic mirror, giving a vacuum focal spot size of 22 μm FWHM and a vacuum peak intensity of 1.3×10^{18} W cm^{-2}.

Methane has a high propensity to form clusters as the gas cools due to adiabatic expansion into the vacuum, and at the backing pressure used of 7 bar (absolute) we estimate an average cluster size of 10^4 molecules (~6 nm radius) and a plasma density of 9×10^{18} cm^{-3}. We used a methane cluster gas as a target, because previous experiments using the same laser parameters have found that the stability and injected charge is improved compared to helium gas jets [22].

The laser temporal profile was controlled using an acousto-optic programmable dispersive filter [23] (Fastlite Dazzler) positioned in the front end of the laser. The Dazzler can adjust both the spectral amplitude and phase, but, to avoid extreme changes to the pulse shape and spectrum affecting the performance of the laser amplifiers,

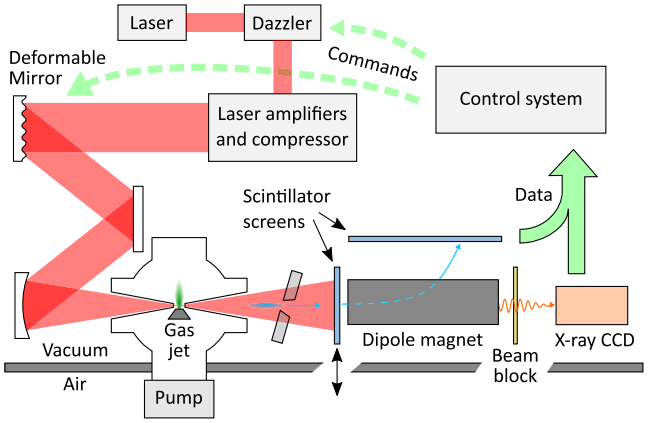

FIG. 1. A schematic diagram of the experimental setup. The laser (light red) is focused into a gas jet (bottom left), accelerating an electron bunch (light blue). A holed glass wedge allows the transmitted laser beam to be diagnosed. Scintillator screens record the electron beam's spatial profile and energy spectrum; the former screen can be removed to improve the resolution of the spectrometer. Betatron x rays are recorded by an x-ray CCD, which is protected from any residual laser light by a thin gold beam block. Data recorded from the scintillator screens and the x-ray CCD are sent to the control system (top right), which is running the optimization algorithm. This sends commands to the Dazzler and the deformable mirror, completing the feedback loop.

or causing damage, we altered only the spectral phase as a polynomial up to fourth order, and we did not manipulate the spectral amplitude at all. This is still enough to control the temporal shape of the laser pulse, particularly stretching it by applying a linear chirp (second-order phase) or skewing it by adding a slow-rising or-falling edge (third-order phase).

The pulse was characterized before the interaction on every shot using frequency-resolved optical gating [24] (Swamp Optics Grenouille). The Dazzler settings that gave the shortest pulse were found by adjusting the settings manually while measuring the pulse length. These settings were used as a starting point when the Dazzler was being controlled by an optimization algorithm. We detected a drift in the pulse shape, amounting to a change in the second-order spectral phase of approximately $+500\ \mathrm{fs}^2$, over the course of about 3000 laser shots (~20 min). We believe this to be the result of heating and thermal expansion of the compressor gratings [25] affecting the first 60 measurements (each averaging 50 shots) at the start of each optimization run before the compressor stabilizes. The algorithm appears to produce good results despite this problem, and indeed it serves as a demonstration that gradual changes in experimental conditions can be accommodated. An important consequence is that care has to be taken to check that the actual spectral phase of the laser pulse matches the phase expected from the applied Dazzler settings.

The spatial profile was controlled by a deformable mirror (DM) with a high-reflectivity dielectric coating located before the parabola. A Shack-Hartmann sensor (Imagine-Optic HASO) was operated in a leakage beam line to monitor the wavefront of the beam. The DM surface was adapted with 55 individual actuators, the effect of which forms a 55-dimensional basis. A more convenient basis can be constructed from the Zernike polynomials, commonly used in the aberration theory [26,27]. In principle, 55 different modes are accessible, but stiffness in the mirror surface limits the amplitudes of the higher-order modes. By retaining only the 18 lowest-order modes (excluding the two tilt modes and the zero-order piston mode), we could reduce the dimensionality of the search space and simplify the optimization task.

The best focus was obtained by optimizing the quality of the focal spot using attenuated laser power. The resulting mirror shape was used throughout the experiment, except when it was the subject of another optimization. This process itself served as a test of the optimization routine. The goal function for this case was

$$G = \sum_{i,j} V_{i,j}^2, \tag{1}$$

where i and j run over the rows and columns in the image of the focal spot and $V_{i,j}$ is the value of the pixel at coordinates (i, j) in the same image. Since $\sum_{i,j} V_{i,j}$ is assumed to be constant (being proportional to the laser energy), this goal function is maximized when the size of the focal spot is minimized. The algorithm used in this case was the Nelder-Mead method (described in Sec. III).

In order to make full use of the repetition rate, the optimization algorithm should be able to run without the need for human intervention. To this end, a control system was created to automatically manage the entire process of evaluating the goal function. This involved the remote configuration of the Dazzler or DM, interfacing with the main laser control system which provided bursts of 50 pulses, acquiring the data, and evaluating the goal function itself. Once this system was operational, it could simply be plugged into the optimization algorithm as any other function would be.

Naturally, a computer system controlling the laser automatically, for an extended period of time without an operator, carries a certain risk of damage to equipment if, for instance, inappropriate laser parameters are selected. The main risk was damage to the laser amplifier chain that might be caused by a choice of Dazzler parameters causing severe spectral clipping. Therefore, conservative limits were placed on the range of Dazzler parameters that could be selected by the algorithm. These limits were enforced in the code responsible for communicating with the Dazzler.

To deal with the high gas load associated with 5 Hz operation, the gas jet was enclosed in a second chamber (an

ISO-160 six-way cross) within the main target chamber. This inner chamber was directly attached to a dedicated large-capacity rotary vacuum pump. It was joined to the main chamber by a pair of 2-mm-diameter pinholes, allowing passage of the laser and electron beam but limiting gas flow. This differential pumping geometry allowed the pressure within the main vacuum chamber to be kept below 3×10^{-4} mbar even when the inner chamber pressure increased above 10^{-2} mbar.

The spectra of the accelerated electron beams were measured using a magnetic spectrometer, consisting of a 0.6 T, 300-mm-long dipole magnet and a Lanex screen. This was calibrated using a magnetic field map, with the calculated energy having a 1σ relative error of 10%. In addition, beam divergence and pointing fluctuations limit the resolution of the spectrometer, with an associated 1σ relative error of 22%. Absolute calibration of the intensity (charge) scale was accomplished by placing a piece of imaging plate in front of (from the electrons' point of view) the Lanex screen and recording a series of electron beams.

The electron beam profile could be measured by inserting a Lanex screen before the entrance to the spectrometer, although this scattered the electrons and degraded the spectrometer's performance, particularly at low energies (below ~20 MeV). The betatron x rays were measured by a deep-depletion CCD (Andor iKon-M) following the magnetic spectrometer, with a filter pack in place to allow the x-ray spectrum to be measured.

III. OPTIMIZATION ALGORITHMS

Two optimization algorithms were tested in this experiment: a differential genetic algorithm and the Nelder-Mead method [28].

Genetic algorithms are stochastic global optimization algorithms. By analogy with the process of natural selection, they maintain a "population" of points in parameter space. The goal function is evaluated at each of these points, and then the most optimal members of the population are selected. These are used to produce a new "generation," replacing the existing population. Over many generations, the algorithm is expected to converge towards the global optimum.

Generally, the new generation is produced by "crossover" and "mutation" operations, both analogous to similarly named genetic processes. The crossover operation randomly combines two points to produce a single point, in such a way that information is taken from both "parents." This is often achieved by copying each parameter from one of the parents, selected at random. The mutation operation is a random perturbation to a single point and is needed to maintain a diverse population and avoid premature convergence to a false optimum.

In our version of the algorithm, the mutation is based on the difference between two randomly selected members of the population, which avoids the need to fine-tune the size of mutations. This idea is taken from a related algorithm, differential evolution [29].

The crossover and mutation operations, and, in particular, the randomness used by them, allow the algorithm to widely explore the parameter space in search of the global optimum. By contrast, the Nelder-Mead method [28] is much more direct. It maintains a simplex, a set of $N + 1$ points in an N-dimensional space, and evaluates the goal function at each point of the simplex. At each step, it attempts to improve the worst point by moving it towards or through the center of the simplex. Although this method is designed for local optimization, it has potential advantages: It is deterministic, and it might converge more quickly than a genetic algorithm, because it moves towards the optimum rather than making random changes.

IV. RESULTS

Using the smallest focus and shortest pulse, electrons were easily and consistently injected and accelerated to energies as high as 100 MeV. The mean charge measured on the electron spectrometer was 55 pC, with a standard deviation of 19 pC. The energy spectrum was continuous, as has been previously observed with cluster gas targets [30], but it also featured a peak of variable size and energy. By averaging over a burst of 50 shots, the quality of the data could be improved enough to be used with the optimization algorithm. The benefit of the averaging can be seen in Figs. 2, which shows electron beam spatial profiles, and 3, which shows electron energy spectra.

We performed a number of different optimization runs. Each started from the shortest pulse, as measured by the Grenouille, and adjusted the Dazzler settings to improve a particular goal function. A selection is presented here. These examples all used a genetic algorithm (as described in Sec. III), with 15 samples per generation, the best four of

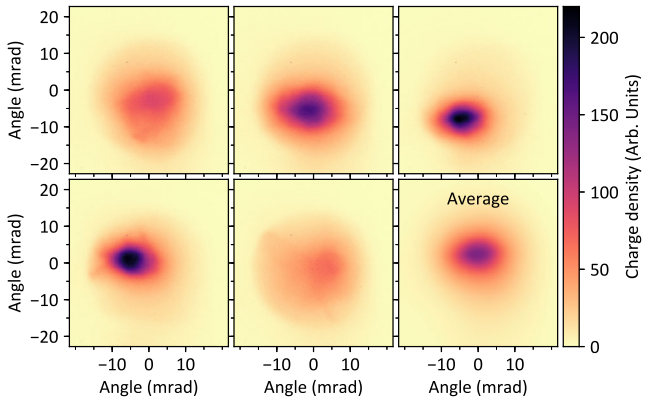

FIG. 2. Five randomly selected images of the electron beam spatial profile, together with a 50-shot average. Each image is independently normalized. The average was taken after overlapping the centroids of the images, to remove the effects of pointing fluctuations. Data were taken using the temporal pulse shape which maximized the accelerated charge.

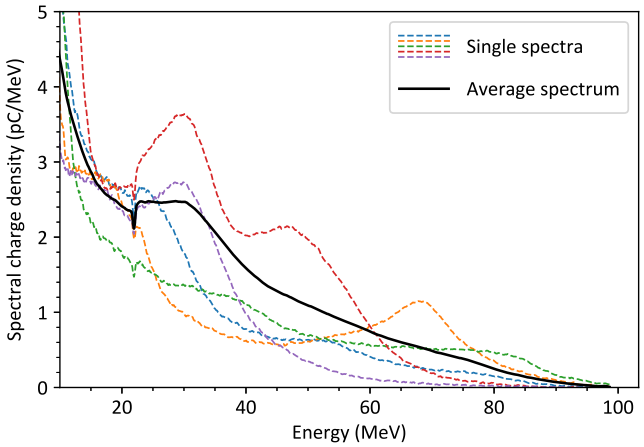

FIG. 3. Five randomly selected electron energy spectra, together with a 50-shot average. The individual spectra correspond to the beam profiles of the same color shown in Fig. 2. Data were taken using the temporal pulse shape which maximized the accelerated charge.

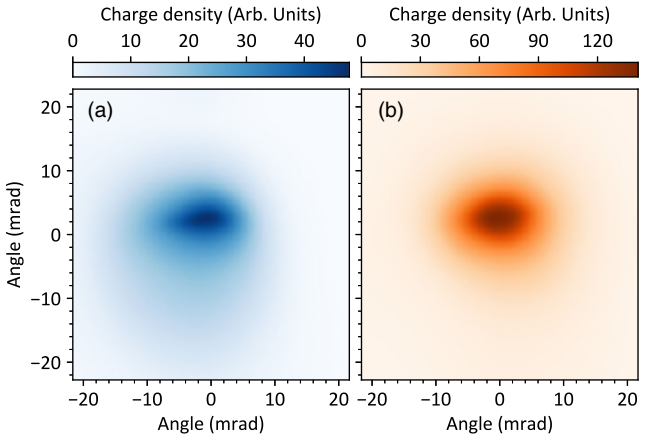

FIG. 4. Images of the electron beam spatial profile (a) before and (b) after optimization. Each image is an average of 50 shots, after removing pointing fluctuations. The color scales are different: The total charge in (b) is 2.1× the total charge in (a).

which were selected to be parents and the best two of which were carried into the next generation unchanged. The number of generations was not chosen beforehand: We terminated the algorithm when it appeared to be making no further improvement.

A. Charge optimization

One of the simplest goal functions is the total amount of accelerated charge. This was quantified by summing the pixel values on the Lanex screen inserted before the magnetic spectrometer. Between the initial conditions (shortest pulse) and the best parameters tested by the algorithm, the total signal measured on the electron profile screen increased by a factor of 2.1. Other measurements of the charge yielded slightly different results: The peak signal increased by a factor of 2.9, while the charge measured on the electron spectrometer increased from 54 to 140 pC, a factor of 2.6 increase, suggesting that a greater fraction of the electrons are being accelerated to high energies. The spatial profiles of the electron beams produced by the shortest and optimized pulses are shown in Figs. 4(a) and 4(b), respectively.

B. Energy spectrum optimization

While more charge is generally desirable, there are many applications where it would be useful to have greater control over the properties of the electron beam. For instance, a tunable x-ray source based on inverse Compton scattering [31] might require a tunable, narrow energy spread electron beam. Although the main methods of controlling the electron energy involve changing the injection location or interaction length, it might also be possible to manipulate the drive laser itself. If so, a feedback system could be used to tune the laser to the conditions needed for a particular value of the electron energy.

As a demonstration, we defined a goal function based on the total charge within a specified energy range. This involved taking an image of the Lanex screen used in the magnetic spectrometer, applying a suitable perspective correction, and then summing the signal within the interval corresponding to the desired energy range. To maximize the spectrometer signal quality, the screen used to measure the electron beam spatial profile was removed.

This is similar to the method used by He *et al.* to control the electron energy distribution [21], although in that case the spatial profile was adapted rather than the temporal. By optimizing the signal within a rectangular mask, He *et al.* were able to control the average energy of a continuous electron spectrum by moving the mask along the dispersion direction. Average energies of 89, 95, and 98 keV were produced in separate runs of the genetic algorithm.

We performed two optimization runs, both using a genetic algorithm to maximize the goal function while varying the spectral phase up to fourth order. Each run used a different target energy range: One run used 29–34 MeV, while the other used 57–68 MeV. The results are shown in Fig. 5(a), which shows, for each run, the initial spectrum, the final spectrum, and the target energy range.

Although both initial spectra were produced by the Dazzler settings appropriate for the shortest pulse, there are significant differences between them, indicating some change in experimental conditions. To account for this difference, we consider the ratio of the optimized spectrum to the initial spectrum, shown in Fig. 5(b).

Both optimization runs resulted in a significant increase in the overall charge: by a factor of 2.3 for the low-energy (blue) case and a factor of 1.6 for the high-energy (orange) case. In the low-energy case this improvement was concentrated around the target interval, as shown in Fig. 5(b),

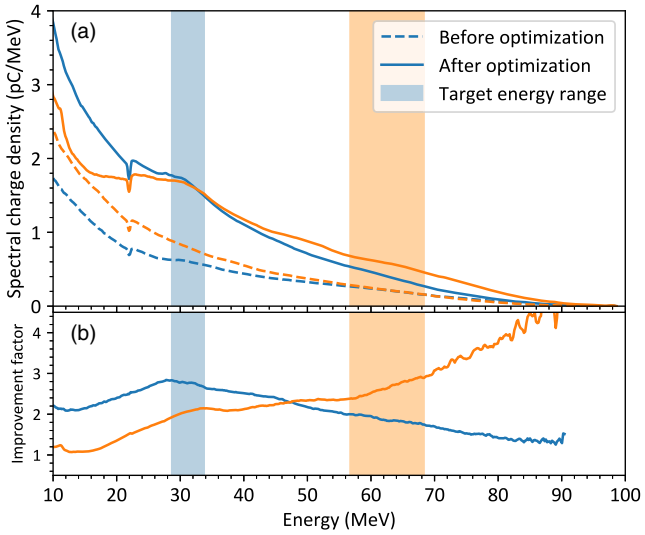

FIG. 5. (a) The electron spectra (averaged over 50 shots) produced at the beginning (dashed lines) and end (solid lines) of two optimization runs. Each optimization run is plotted in a different color, and each was targeted at a different energy interval (shaded regions). (b) The improvement factor: the ratio of the final spectrum to the initial spectrum.

while in the high-energy case the most improvement occurs at higher energies, above 50 MeV. The difference between the two optimized spectra results in different average energies: 23.0 and 27.5 MeV for the low-energy and high-energy optimization runs, respectively, although these figures do not include data below the spectrometer's low-energy cutoff at 5 MeV. Before optimization, the average energies were 23.8 and 21.7 MeV, respectively.

The conditions used here generated continuous energy spectra, and changes to the shape achieved by the algorithm were rather limited. By repeating the experiment in a more controlled LWFA regime, as shown in previous experiments with the same laser [3], we may be able to produce narrow energy spread beams fully tunable through pulse shaping.

C. Optimization using a deformable mirror

As well as using the Dazzler to temporally shape the laser pulse, we were able to employ the DM to control its spatial profile. For the runs described so far, the mirror was adjusted (using the Nelder-Mead method) to minimize the size of the focal spot as observed by a camera. However, an accurate measurement of the focal spot of a high-power laser is disruptive, because it requires orders of magnitude of attenuation. An advantage of the direct optimization of experimental parameters using the feedback loop is that, if the DM is initiated in a reasonable starting position, the ideal focal spot can be achieved without needing to measure it directly. Furthermore, in some cases, it may be desirable to create a nonperfect focus, for example, to

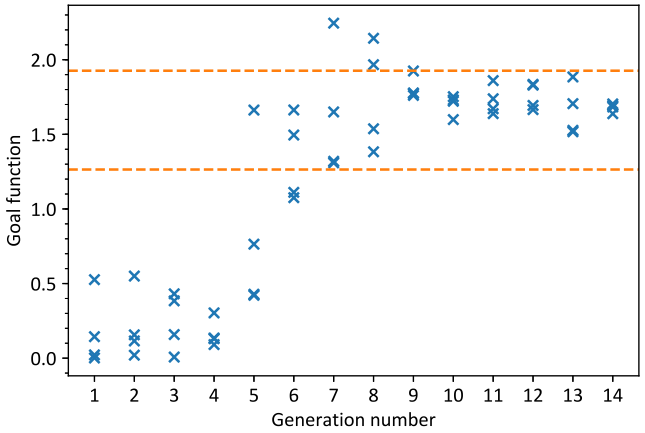

FIG. 6. The progress of optimization using the deformable mirror, described in more detail in the text. Only the best four measurements per generation are shown; these are the ones used to form the next generation. The dashed lines indicate two estimates of the goal function that would be achieved under similar conditions with the best focus.

augment electron oscillations to increase the energy of betatron radiation [32].

To investigate this, we carried out an optimization run with initial DM settings at the midpoint of the voltage range, leading to some aberrations in the focal spot. The goal function was based on an image of the electron profile Lanex screen:

$$G = \sum_{i,j} V_{i,j}^{1.5}, \tag{2}$$

where i and j run over the rows and columns in the image and $V_{i,j}$ is the value of the pixel at coordinates (i, j). The sum was taken only over the pixels containing the screen. This optimizes charge with the nonlinearity intended to favor beams with less divergence. The values of the goal function as the optimization progressed are shown in Fig. 6. As expected, the LWFA performance is poor in generation 1 with the unmodified focal spot. A dramatic improvement occurs in generations 5–7, reaching levels comparable with the values of 1.25 and 1.9 (marked in Fig. 6 with dashed lines) measured during the previous run conducted using the best focal spot and an equivalent temporal profile. This suggests that the beam evolved to a high-quality focus well suited for driving the LWFA.

A relatively large number of generations were needed for the optimum to be reached, perhaps because of the increased dimensionality of the parameter space—18 dimensions rather than three. While the problem can be alleviated by the use of the Zernike polynomials as a basis, the interesting region still forms a much smaller fraction of the parameter space's total hypervolume and so is much harder to find. Data analysis in high-dimensional spaces involves many problems like this, a phenomenon which has been called the "curse of dimensionality" [33]. This is a

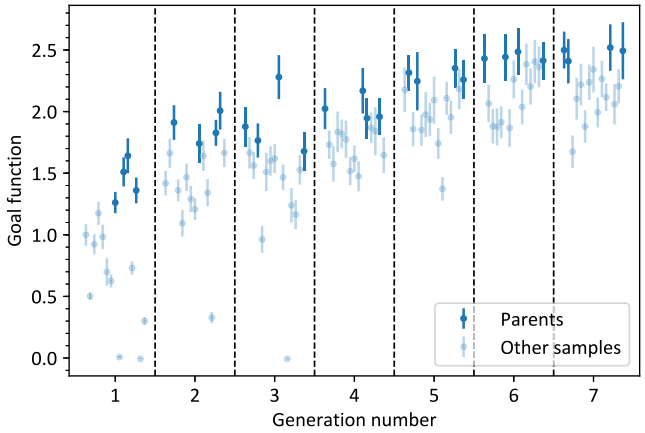

FIG. 7. Every measurement of the goal function during a single optimization run. The four best samples from each generation, used as parents for the next, are highlighted. The vertical axis is normalized to the first sample, taken with a fully compressed laser pulse.

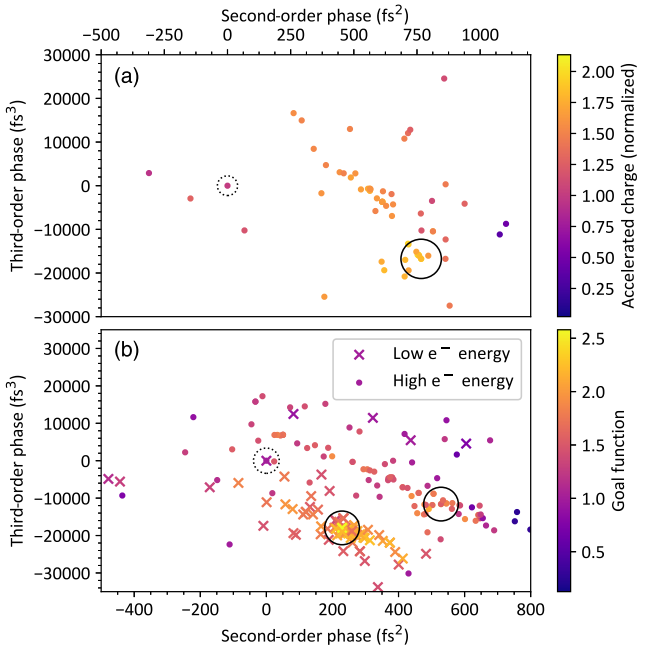

FIG. 8. Dazzler settings (second- and third-order phases only; the fourth-order phase is not shown) used during optimization. In both cases, the initial settings, giving the best pulse compression, are indicated by a dotted circle, while the optimal settings are indicated by a solid circle. (a) Charge optimization. The amount of accelerated charge is indicated by the color of the marker. (b) Energy spectrum optimization, as shown in Fig. 5. The value of the goal function (see the text) is indicated by the color of the marker. The shape of the marker indicates the target energy range. Some points have been omitted for legibility.

difficult problem to overcome, but better results might be achieved in this case if the settings of the genetic algorithm were changed.

V. DISCUSSION

Although optimization is a useful practical technique by itself, learning from the results requires us to understand what the algorithm is doing and why it is doing it. First, we can plot the result of each evaluation of the goal function, as shown in Fig. 7. The data for this plot were taken from the low-energy electron energy spectrum optimization run, as shown in Fig. 5 and described above. The main feature is a gradual increase of the goal function from one generation to the next until it eventually saturates. The fact that some individual measurements are worse than preceding ones is expected behavior: As the improvements are made by random changes to existing parameters, some of these changes will cause the goal function to get worse. The algorithm promises only to improve the maximum over time.

Figure 8 shows the Dazzler settings selected by the algorithm during the three optimization runs discussed. The parameter space was three-dimensional, consisting of the second, third, and fourth polynomial orders of the spectral phase. Only small changes to the fourth order were applied, and this is not shown here. Each point denotes a single measurement, with its color indicating the measured value of the goal function and its position indicating the second- and third-order phase terms applied.

Figure 8(a) shows the measurements made during the optimization of the accelerated charge. The initial conditions (a fully compressed pulse) are indicated by a dotted circle, while the optimal conditions found by the algorithm are indicated by a solid circle. Several groups have found that adding a positive second-order spectral phase (linear chirp) improves the performance of a laser wakefield accelerator [35–38]. In this case, varying the third-order phase as well improves the results further, increasing the accelerated charge by approximately 20%.

Figure 8(b) shows the measurements made during the two optimization runs based on the electron energy spectrum, the final results of which were shown in Fig. 5. Some data have been omitted so that the graph remains legible. Again, the optimal settings differ from the initial settings in both second- and third-order phases. More interestingly, the optimal settings for the two optimization runs are clearly separated, providing further evidence that the difference between the spectra in Fig. 5 is a real effect caused by applying different Dazzler settings.

These observations are broadly consistent with the results of Kim *et al.*, who found that applying a positive second-order phase and a negative third-order phase increased the performance of a wakefield accelerator driven by a petawatt laser [38,39]. The key advance made in this work is that both parameters are optimized simultaneously, which is necessary when the optimal value of the second-order phase at zero third-order phase differs from the globally optimal value. The slanted shape of the data in both panels in Fig. 8 suggests that this is the case here.

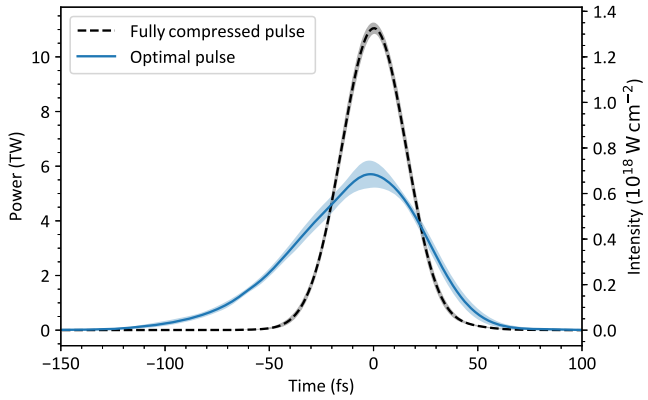

FIG. 9. The temporal shapes of the initial (fully compressed) and optimal pulses from the charge optimization run. The optimal pulse is approximately twice as long and has a pronounced asymmetry. The shaded region indicates the root mean square variation from shot to shot.

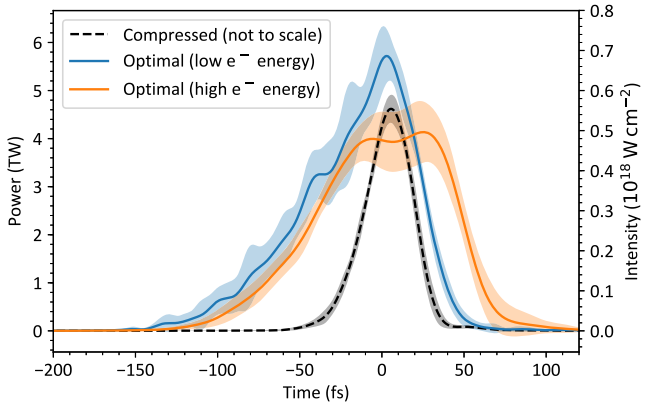

FIG. 10. The temporal shapes of two pulses optimized for different electron energy ranges, as well as the fully compressed pulse used as the start point of the optimization runs.

The conditions in Kim *et al.* were significantly different and do not exhibit the same problem (see Fig. 4 in Ref. [39]).

A. Pulse shapes

We gain insight into the physical processes responsible for the improvements in signal by considering the pulse shapes measured by the Grenouille. Figure 9 shows the pulse produced by optimizing the accelerated charge, and Fig. 10 shows the two pulses that increased charge and optimized the electron energy spectra in two different ranges. The fully compressed pulses used as the starting points are also displayed. In all cases, the optimized pulse has a longer duration and is skewed, having a slow-rising edge. This is in agreement with the Dazzler parameters (Fig. 8), which show an additional positive second-order (mainly affecting pulse length) and negative third-order (mainly affecting skewness) spectral phase. The pulse optimized for a higher electron energy has a flatter profile, reducing the peak intensity but maintaining it for longer.

Despite operating in a very different regime, the optimal pulse shape is similar to the one found by Kim *et al.*, having a slow-rising edge and a fast-falling edge. Their conclusion, supported by particle-in-cell (PIC) simulations, was that this shape significantly reduced the amount of self-modulation suffered by the laser pulse [38,39]. This led to a more stable acceleration process, maintained over a longer distance. The same mechanism may be enhancing the accelerator performance under our conditions. Of particular interest is the slight flattening of the laser pulse when it is optimized for a higher electron energy in Fig. 10. This suggests that higher intensities are not necessarily desirable for reaching high electron energies, perhaps because nonlinear effects cause more self-modulation of the laser pulse.

Another possible explanation is direct laser acceleration, which has been observed when the laser pulse is long enough to overlap the accelerating phase of the wakefield [40]. In that case, the length and shape of the pulse would affect the acceleration process more directly.

It is also important to take account of the presence of solid-density clusters in the gas. Ion expansion determines the extent of electron extraction from the cluster cores and the heating of these electrons. These processes occur on the timescale of the laser pulse and are strongly affected by the temporal intensity profile. The slower-rising edge found here is known to increase the electron temperature [41], and this could lead to an increase in the rate of self-injection [42].

Ideally, we would perform PIC simulations of the laser-plasma interaction using shaped intensity profiles as the drive pulse. However, the clusters make this extremely difficult, because they introduce physics on a length scale much smaller than the laser wavelength. We cannot neglect the clusters, as it has previously been shown that they substantially increase the amount of charge accelerated [22,30,43]. As such, these simulations are beyond the scope of this work.

VI. COMPARISON OF OPTIMIZATION ALGORITHMS

Previous active feedback experiments [19,20] have used genetic algorithms to solve the optimization problem. These are able to solve complex problems, but they require a certain amount of fine-tuning. The fact that they make extensive use of random numbers may also be a problem, as performing the same optimization twice under the same conditions will lead to different results. Even if the sequence of random numbers is repeatable, slight changes in experimental measurements can mean that they are used in a different order. Given these drawbacks, it is worth investigating other algorithms to see whether they are any better.

With this in mind, we performed some optimization runs using the Nelder-Mead method, described in Sec. III.

TABLE I. Two comparisons between the genetic algorithm (GA) and Nelder-Mead method (NM). Two criteria are shown: the improvement in the goal function (shown as a percentage) and the number of function evaluations required to attain it. In each case, the better result is shown in bold.

Goal function	Charge		X rays	
Algorithm	GA	NM	GA	NM
Improvement	**130%**	95%	70%	**100%**
Evaluations	45	**25**	75	**40**

The main criteria that determine which algorithm to use are the final value of the goal function and the number of function evaluations required to attain it. These depend not only on the choice of algorithm, but also on any settings used by the algorithm, as well as the precise problem being solved. There is also a random component, due to experimental errors.

Although the genetic algorithm, being designed to find the global optimum, is expected to produce a better result in many cases, this may come at the cost of requiring a large number of function evaluations. Having a more optimized electron beam is not necessarily beneficial if there is less time to use it. The Nelder-Mead method has the advantage that it makes no explicit use of random numbers—experimental errors may be considered a source of random numbers that are used implicitly by the algorithm. This might make its results more consistent and reproducible.

Unfortunately, a quantitative comparison of the two algorithms is difficult. It would require many runs and many tens of thousands of laser shots to average over random fluctuations, and this would ideally be repeated for many different goal functions. This was not possible during the timeframe of this experiment. In two cases, we performed the same optimization twice with both algorithms: One involved maximizing the total x-ray signal, and the other involved maximizing the charge—specifically, the goal function shown in Eq. (2). In each case, the sequence of measurements was analyzed to give the improvement in the goal function and the number of function evaluations required. The results are shown in Table I.

In the charge optimization case, the genetic algorithm produced the better result, whereas the Nelder-Mead method performed better in the x-ray optimization case. In both cases, the Nelder-Mead method reached the optimum faster. Although this is not enough data to draw any firm conclusions, it does at least suggest that both algorithms are viable.

This is significant, because global optimization algorithms, including genetic algorithms, are designed to avoid becoming trapped in a local maximum. This capability is crucial for solving many problems but comes at the cost of efficiency. The performance of the Nelder-Mead method in this case suggests that local optimization algorithms work well in this case, opening up a much larger space of possible algorithms to explore.

The genetic algorithm has another significant advantage: its lack of "long-term memory." Measurements are only ever compared within a single generation, which naturally accommodates gradual changes in experimental conditions. By contrast, the Nelder-Mead method retains measurements indefinitely. Changes in experimental conditions could result in a previously optimal point in parameter space becoming worse. The algorithm would retain the high value of the goal function previously measured, causing it to make poor decisions. In principle, this problem could be mitigated by periodically discarding cached measurements and retaking them, although this would require more function evaluations.

It is not clear whether the observed drift in the spectral phase affected our work with the Nelder-Mead algorithm. It is interesting to note that the charge optimization run, when the Nelder-Mead algorithm outperformed the genetic algorithm, took place immediately after a run lasting 50 min. This might have allowed the system to achieve thermal stability. By contrast, the x-ray optimization run, when the Nelder-Mead algorithm did not perform so well, followed a break of nearly 25 min. This would have provided enough time for the system to cool down. This is consistent with the hypothesis that the Nelder-Mead method performs poorly if experimental conditions change during the run, but we cannot say any more than that without additional data.

VII. CONCLUSIONS AND FURTHER WORK

We have shown that optimization using active feedback can be applied to a laser wakefield accelerator using a higher laser energy (450 mJ) and lower repetition rates (5 Hz) than previous work [20]. This scale of laser is already sufficient to produce 200 MeV electrons [44] and hard x-ray beams up to the 1 MeV level [12,45].

Extending the same principles to high repetition rate petawatt-class lasers would be an obvious next step offering high-power laser-driven secondary sources with additional control and better stability for applications.

The benefit of active feedback is that performance can be improved without the need to predict the ideal pulse shape. Enhancements can result from the evolution of the pulse to a complicated spatiotemporal focal profile. The underlying physics can then be fully understood by feeding this information into theoretical models.

Starting from a fully temporally compressed drive pulse, our genetic algorithm was able to more than double the amount of accelerated charge by altering the temporal profile of the pulse to be longer and have a slow-rising edge. This shape may have improved the performance by reducing laser self-modulation and by tailoring the expansion of the clusters to optimize injection into the wakefield.

By optimizing for the charge measured within a given energy range, we were able to produce two different pulse

shapes that resulted in two different averaged electron energy spectra: one with an average energy of 23 MeV and another with an average energy of 27.5 MeV. Both pulses were able to accelerate significantly more charge (by a factor of 1.6–2.3) than the fully compressed pulse, and both had the same long duration and slow-rising edge. The main difference was in the peak, which was longer and flatter in the high-energy case. This may indicate optimization for a longer depletion length.

Although previous work focused on genetic algorithms, many other optimization algorithms have been developed over the years for many different purposes. As an example, we tested the Nelder-Mead method, which is not designed for global optimization but has its own advantages, notably being more deterministic, and might be useful either by itself or together with a global optimization algorithm.

However, it is also possible that neither of these algorithms are well suited to the task of active feedback in a laser wakefield accelerator. Evaluation of the goal function both is noisy (meaning it includes a random error term) and consumes valuable experimental time, and specialized algorithms may be needed to cope with this. One example is a robust conjugate direction search [46], which has been successfully applied to perform online optimization at a number of conventional accelerator facilities. Another is Bayesian optimization [47], which uses statistical models to determine the best points at which to sample the goal function. Another potentially interesting approach is extremum seeking [48], a well-established control technique that is used in many fields of engineering. The application of these and other algorithms to the LWFA would be a useful subject of further study.

The data and code used to generate the plots in this paper are available on Zenodo at [49].

ACKNOWLEDGMENTS

The authors thank the staff of the Central Laser Facility for assistance with the experiment. We acknowledge funding from Science and Technology Facilities Council Grant No. ST/J002062/1 and the European Network for Novel Accelerators, Newton-Bhabha, and Helmholtz Accelerator Initiative programs. S. E. is supported by an Atomic Weapons Establishment Collaborative Awards in Science and Engineering award. J. H., J. A. N., K. K., and A. G. R. T. acknowledge funding from National Science Foundation Grant No. 1535628 and Department of Energy Grant No. DE-SC0016804.

[1] T. Tajima and J. M. Dawson, Laser Electron Accelerator, Phys. Rev. Lett. **43**, 267 (1979).

[2] J. M. Dawson, Nonlinear electron oscillations in a cold plasma, Phys. Rev. **113**, 383 (1959).

[3] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas, J. L. Collier, A. E. Dangor, E. J. Divall, P. S. Foster, J. G. Gallacher, C. J. Hooker, D. A. Jaroszynski, A. J. Langley, W. B. Mori, P. A. Norreys, F. S. Tsung, R. Viskup, B. R. Walton, and K. Krushelnick, Monoenergetic beams of relativistic electrons from intense laser–plasma interactions, Nature (London) **431**, 535 (2004).

[4] C. G. R. Geddes, C. Toth, J. Van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans, High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding, Nature (London) **431**, 538 (2004).

[5] J. Faure, Y. Glinec, A. Pukhov, S. Kiselev, S. Gordienko, E. Lefebvre, J.-P. Rousseau, F. Burgy, and V. Malka, A laser–plasma accelerator producing monoenergetic electron beams, Nature (London) **431**, 541 (2004).

[6] X. Wang *et al.*, Quasi-monoenergetic laser-plasma acceleration of electrons to 2 GeV, Nat. Commun. **4**, 1988 (2013).

[7] W. P. Leemans, A. J. Gonsalves, H.-S. Mao, K. Nakamura, C. Benedetti, C. B. Schroeder, C. Tóth, J. Daniels, D. E. Mittelberger, S. S. Bulanov, J.-L. Vay, C. G. R. Geddes, and E. Esarey, Multi-GeV Electron Beams from Capillary-Discharge-Guided Subpetawatt Laser Pulses in the Self-Trapping Regime, Phys. Rev. Lett. **113**, 245002 (2014).

[8] M. Fuchs, R. Weingartner, A. Popp, Z. Major, S. Becker, J. Osterhoff, I. Cortrie, B. Zeitler, R. Horlein, G. D. Tsakiris, U. Schramm, T. P. Rowlands-Rees, S. M. Hooker, D. Habs, F. Krausz, S. Karsch, and F. Gruner, Laser-driven soft-X-ray undulator source, Nat. Phys. **5**, 826 (2009).

[9] H.-P. Schlenvoigt, K. Haupt, A. Debus, F. Budde, O. Jackel, S. Pfotenhauer, H. Schwoerer, E. Rohwer, J. G. Gallacher, E. Brunetti, R. P. Shanks, S. M. Wiggins, and D. A. Jaroszynski, A compact synchrotron radiation source driven by a laser-plasma wakefield accelerator, Nat. Phys. **4**, 130 (2008).

[10] A. Rousse, K. T. Phuoc, R. Shah, A. Pukhov, E. Lefebvre, V. Malka, S. Kiselev, F. Burgy, J.-P. Rousseau, D. Umstadter, and D. Hulin, Production of a keV X-Ray Beam from Synchrotron Radiation in Relativistic Laser-Plasma Interaction, Phys. Rev. Lett. **93**, 135005 (2004).

[11] S. Kneip, S. R. Nagel, C. Bellei, N. Bourgeois, A. E. Dangor, A. Gopal, R. Heathcote, S. P. D. Mangles, J. R. Marquès, A. Maksimchuk, P. M. Nilson, K. T. Phuoc, S. Reed, M. Tzoufras, F. S. Tsung, L. Willingale, W. B. Mori, A. Rousse, K. Krushelnick, and Z. Najmudin, Observation of Synchrotron Radiation from Electrons Accelerated in a Petawatt-Laser-Generated Plasma Cavity, Phys. Rev. Lett. **100**, 105006 (2008).

[12] S. Corde, K. Ta Phuoc, G. Lambert, R. Fitour, V. Malka, A. Rousse, A. Beck, and E. Lefebvre, Femtosecond x rays from laser-plasma accelerators, Rev. Mod. Phys. **85**, 1 (2013).

[13] S. Fourmaux, S. Corde, K. T. Phuoc, P. Lassonde, G. Lebrun, S. Payeur, F. Martin, S. Sebban, V. Malka, A. Rousse, and J. C. Kieffer, Single shot phase contrast imaging using laser-produced Betatron x-ray beams, Opt. Lett. **36**, 2426 (2011).

[14] S. Kneip, C. McGuffey, F. Dollar, M. S. Bloom, V. Chvykov, G. Kalintchenko, K. Krushelnick, A. Maksimchuk, S. P. D. Mangles, T. Matsuoka, Z. Najmudin, C. A. J.

Palmer, J. Schreiber, W. Schumaker, A. G. R. Thomas, and V. Yanovsky, X-ray phase contrast imaging of biological specimens with femtosecond pulses of betatron radiation from a compact laser plasma wakefield accelerator, Appl. Phys. Lett. **99**, 093701 (2011).

[15] Z. Najmudin *et al.*, Compact laser accelerators for X-ray phase-contrast imaging, Phil. Trans. R. Soc. A **372**, 20130032 (2014).

[16] P. Mason, M. Divoký, K. Ertel, J. Pilař, T. Butcher, M. Hanuš, S. Banerjee, J. Phillips, J. Smith, M. D. Vido, A. Lucianetti, C. Hernandez-Gomez, C. Edwards, T. Mocek, and J. Collier, Kilowatt average power 100 J-level diode pumped solid state laser, Optica **4**, 438 (2017).

[17] C. L. Haefner *et al.*, High average power, diode pumped petawatt laser systems: A new generation of lasers enabling precision science and commercial applications, Proc. SPIE Int. Soc. Opt. Eng. **10241**, 1024102 (2017).

[18] H. Rabitz, R. d. Vivie-Riedle, M. Motzkus, and K. Kompa, Whither the future of controlling quantum phenomena?, Science **288**, 824 (2000).

[19] S. Sharma, H. Singh, and G. G. Balint-Kurti, Genetic algorithm optimization of laser pulses for molecular quantum state excitation, J. Chem. Phys. **132**, 064108 (2010).

[20] Z.-H. He, B. Hou, V. Lebailly, J. A. Nees, K. Krushelnick, and A. G. R. Thomas, Coherent control of plasma dynamics, Nat. Commun. **6**, 7156 (2015).

[21] Z.-H. He, B. Hou, G. Gao, V. Lebailly, J. A. Nees, R. Clarke, K. Krushelnick, and A. G. R. Thomas, Coherent control of plasma dynamics by feedback-optimized wavefront manipulation, Phys. Plasmas **22**, 056704 (2015).

[22] S. V. Rozario *et al.*, Investigating the effects of clustering in a laser wakefield accelerator, CLF annual report, 2015–2016, https://www.clf.stfc.ac.uk/Pages/20%20-%20Rozario.pdf.

[23] P. Tournois, Acousto-optic programmable dispersive filter for adaptive compensation of group delay time dispersion in laser systems, Opt. Commun. **140**, 245 (1997).

[24] D. J. Kane and R. Trebino, Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating, IEEE J. Quantum Electron. **29**, 571 (1993).

[25] V. Leroux, S. W. Jolly, M. Schnepp, T. Eichner, S. Jalas, M. Kirchen, P. Messner, C. Werle, P. Winkler, and A. R. Maier, Wavefront degradation of a 200 TW laser from heat-induced deformation of in-vacuum compressor gratings, Opt. Express **26**, 13061 (2018).

[26] R. J. Noll, Zernike polynomials and atmospheric turbulence, J. Opt. Soc. Am. **66**, 207 (1976).

[27] L. Zhu, P.-C. Sun, D.-U. Bartsch, W. R. Freeman, and Y. Fainman, Wave-front generation of Zernike polynomial modes with a micromachined membrane deformable mirror, Appl. Opt. **38**, 6019 (1999).

[28] J. A. Nelder and R. Mead, A Simplex Method for Function Minimization, Comput. J. **7**, 308 (1965).

[29] R. Storn and K. Price, Differential evolution – A simple and efficient heuristic for global optimization over continuous spaces, J. Global Optim. **11**, 341 (1997).

[30] M. Mirzaie, N. A. M. Hafz, S. Li, K. Gao, G. Li, Q. ul Ain, and J. Zhang, Laser acceleration in argon clusters and gas media, Plasma Phys. Controlled Fusion **58**, 034014 (2016).

[31] K. Khrennikov, J. Wenz, A. Buck, J. Xu, M. Heigoldt, L. Veisz, and S. Karsch, Tunable All-Optical Quasimonochromatic Thomson X-Ray Source in the Nonlinear Regime, Phys. Rev. Lett. **114**, 195003 (2015).

[32] S. P. D. Mangles, G. Genoud, S. Kneip, M. Burza, K. Cassou, B. Cros, N. P. Dover, C. Kamperidis, Z. Najmudin, A. Persson, J. Schreiber, F. Wojda, and C.-G. Wahlström, Controlling the spectrum of x-rays generated in a laser-plasma accelerator by tailoring the laser wavefront, Appl. Phys. Lett. **95**, 181106 (2009).

[33] This term was apparently coined by Richard Belman in 1961 to refer to the challenges of optimizing a function by exhaustively searching its domain. Similar problems exist in other fields, including statistics [34].

[34] G. V. Trunk, A problem of dimensionality: A simple example, IEEE Trans. Pattern Anal. Mach. Intell. **PAMI-1**, 306 (1979).

[35] B. S. Rao, A. Moorti, P. A. Naik, and P. D. Gupta, Effect of chirp on self-modulation and laser wakefield electron acceleration in the regime of quasimonoenergetic electron beam generation, Phys. Rev. ST Accel. Beams **16**, 091301 (2013).

[36] C. Liu, J. Zhang, S. Chen, G. Golovin, S. Banerjee, B. Zhao, N. Powers, I. Ghebregziabher, and D. Umstadter, Adaptive-feedback spectral-phase control for interactions with transform-limited ultrashort high-power laser pulses, Opt. Lett. **39**, 80 (2014).

[37] T. Z. Zhao, K. Behm, C. F. Dong, X. Davoine, S. Y. Kalmykov, V. Petrov, V. Chvykov, P. Cummings, B. Hou, A. Maksimchuk, J. A. Nees, V. Yanovsky, A. G. R. Thomas, and K. Krushelnick, High-Flux Femtosecond X-Ray Emission from Controlled Generation of Annular Electron Beams in a Laser Wakefield Accelerator, Phys. Rev. Lett. **117**, 094801 (2016).

[38] H. T. Kim, V. B. Pathak, K. Hong Pae, A. Lifschitz, F. Sylla, J. H. Shin, C. Hojbota, S. K. Lee, J. H. Sung, H. W. Lee, E. Guillaume, C. Thaury, K. Nakajima, J. Vieira, L. O. Silva, V. Malka, and C. H. Nam, Stable multi-GeV electron accelerator driven by waveform-controlled PW laser pulses, Sci. Rep. **7**, 10203 (2017).

[39] J. Shin, H. T. Kim, V. B. Pathak, C. Hojbota, S. K. Lee, J. H. Sung, H. W. Lee, J. W. Yoon, C. Jeon, K. Nakajima, F. Sylla, A. Lifschitz, E. Guillaume, C. Thaury, V. Malka, and C. H. Nam, Quasi-monoenergetic multi-GeV electron acceleration by optimizing the spatial and spectral phases of PW laser pulses, Plasma Phys. Controlled Fusion **60**, 064007 (2018).

[40] J. L. Shaw, N. Lemos, L. D. Amorim, N. Vafaei-Najafabadi, K. A. Marsh, F. S. Tsung, W. B. Mori, and C. Joshi, Role of Direct Laser Acceleration of Electrons in a Laser Wakefield Accelerator with Ionization Injection, Phys. Rev. Lett. **118**, 064801 (2017).

[41] M. J. V. Streeter *et al.*, Temporal feedback control of high-intensity laser pulses to optimize ultrafast heating of atomic clusters, Appl. Phys. Lett. **112**, 244101 (2018).

[42] C. B. Schroeder, E. Esarey, B. A. Shadwick, and W. P. Leemans, Trapping, dark current, and wave breaking in nonlinear plasma waves, Phys. Plasmas **13**, 033103 (2006).

[43] Y. Fukuda, Y. Akahane, M. Aoyama, Y. Hayashi, T. Homma, N. Inoue, M. Kando, S. Kanazawa, H. Kiriyama, S. Kondo, H. Kotaki, S. Masuda, M. Mori, A. Yamazaki, K. Yamakawa, E. Y. Echkina, I. N. Inovenkov, J. Koga, and

S. V. Bulanov, Ultrarelativistic electron generation during the intense, ultrashort laser pulse interaction with clusters, Phys. Lett. A **363**, 130 (2007).

[44] C. G. Geddes, S. Rykovanov, N. H. Matlis, S. Steinke, J.-L. Vay, E. H. Esarey, B. Ludewigt, K. Nakamura, B. J. Quiter, C. B. Schroeder, C. Toth, and W. P. Leemans, Compact quasi-monoenergetic photon sources from laser-plasma accelerators for nuclear detection and characterization, Nucl. Instrum. Methods Phys. Res., Sect. B **350**, 116 (2015).

[45] F. Albert and A. G. R. Thomas, Applications of laser wakefield accelerator-based light sources, Plasma Phys. Controlled Fusion **58**, 103001 (2016).

[46] X. Huang, J. Corbett, J. Safranek, and J. Wu, An algorithm for online optimization of accelerators, Nucl. Instrum. Methods Phys. Res., Sect. A **726**, 77 (2013).

[47] E. Brochu, V. M. Cora, and N. de Freitas, A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning, arXiv:1012.2599.

[48] E. Schuster, C. Xu, N. Torres, E. Morinaga, C. Allen, and M. Krstic, Beam matching adaptive control via extremum seeking, Nucl. Instrum. Methods Phys. Res., Sect. A **581**, 799 (2007).

[49] https://doi.org/10.5281/zenodo.2641789.