

Demonstration of electron acceleration in a laser-driven dielectric microstructure

E. A. Peralta¹, K. Soong¹, R. J. England², E. R. Colby², Z. Wu², B. Montazeri³, C. McGuinness¹, J. McNeur⁴, K. J. Leedle³, D. Walz², E. B. Sozer⁴, B. Cowan⁵, B. Schwartz⁵, G. Travish⁴ & R. L. Byer¹

The enormous size and cost of current state-of-the-art accelerators based on conventional radio-frequency technology has spawned great interest in the development of new acceleration concepts that are more compact and economical. Micro-fabricated dielectric laser accelerators (DLAs) are an attractive approach, because such dielectric microstructures can support accelerating fields one to two orders of magnitude higher than can radio-frequency cavity-based accelerators. DLAs use commercial lasers as a power source, which are smaller and less expensive than the radio-frequency klystrons that power today's accelerators. In addition, DLAs are fabricated via low-cost, lithographic techniques that can be used for mass production. However, despite several DLA structures being proposed recently^{1–4}, no successful demonstration of acceleration in these structures has so far been shown. Here we report high-gradient (beyond 250 MeV m⁻¹) acceleration of electrons in a DLA. Relativistic (60-MeV) electrons are energy-modulated over 563 ± 104 optical periods of a fused silica grating structure, powered by a 800-nm-wavelength mode-locked Ti:sapphire laser. The observed results are in agreement with analytical models and electrodynamic simulations. By comparison, conventional modern linear accelerators operate at gradients of 10–30 MeV m⁻¹, and the first linear radio-frequency cavity accelerator was ten radio-frequency periods (one metre) long with a gradient of approximately 1.6 MeV m⁻¹ (ref. 5). Our results set the stage for the development of future multi-staged DLA devices composed of integrated on-chip systems. This would enable compact table-top accelerators on the MeV–GeV (10⁶–10⁹ eV) scale for security scanners and medical therapy, university-scale X-ray light sources for biological and materials research, and portable medical imaging devices, and would substantially reduce the size and cost of a future collider on the multi-TeV (10¹² eV) scale.

In a dielectric laser accelerator, the large-amplitude fields responsible for the accelerating force are provided by a laser pulse. Of particular interest are lasers that operate in the optical or near-infrared (NIR) portion of the electromagnetic spectrum to take advantage of the low loss and high damage threshold of dielectric materials at these wavelengths. Direct acceleration of electrons using a NIR laser was demonstrated using inverse transition radiation⁶. That approach, however, requires a material boundary in the path of the electron beam and has a limited interaction length. To achieve scalable acceleration, the electric field must have a component parallel to the electron beam trajectory, and a (speed-of-light) phase velocity that matches that of the (relativistic) particle beam.

Previous efforts to accelerate particles with lasers have used a variety of techniques to phase-match the electron beam to a co-propagating laser field. The inverse Cerenkov accelerator^{7,8} uses a gas to slow the phase velocity of the laser light. The inverse free electron laser accelerator^{9,10} requires an undulator at its resonance condition¹¹ to give the electron a synchronous velocity component. However, using a gas cell leads to distortions of the electron beam due to multiple scattering, and use of a

permanent magnet undulator limits the compactness of such a structure and introduces deleterious synchrotron radiation effects.

An alternative way to satisfy the phase-velocity condition is by creating tailored longitudinal modes in near-field structures. One such approach, the inverse Smith–Purcell accelerator¹², has been demonstrated using a metallic grating¹³, and more recently, in a dielectric grating with a NIR laser¹⁴. In configurations that minimize the transverse forces, these open structures do not support speed-of-light longitudinal eigenmodes and are therefore useful only at sub-relativistic particle energies. Additionally, they produce an exponentially decaying accelerating field pattern, which distorts the beam. To address these issues, designs using waveguides and photonic crystals have been proposed^{1–3}. These structures present challenging fabrication tolerances, and the required modes are difficult to excite efficiently; as a result, no successful demonstration has been shown. To circumvent these challenges, we use a planar, phase-reset grating accelerator structure⁴, which permits a relatively simple fabrication process and excitation mechanism.

The DLA is fabricated out of fused silica¹⁵ and operates at the wavelength ($\lambda = 800$ nm) of commercially available Ti:sapphire lasers. The structure consists of two opposing binary gratings of period λ_p , separated by a vacuum gap of height g , where the electron beam travels perpendicular to the grating rulings. A scanning electron microscope (SEM) image of the structure's longitudinal cross-section is shown in Fig. 1a. To generate the required accelerating fields, a transverse magnetic (TM) mode laser pulse is incident on the structure perpendicular to both the electron beam direction of propagation and the plane of the gratings. The laser pulse generates a series of grating diffraction modes inside the vacuum channel, and by matching the grating period to the laser wavelength ($\lambda_p = \lambda$) we achieve phase synchronicity between the strong first space harmonic and the particle beam¹⁶.

The structure essentially acts as a longitudinally periodic phase mask, where each grating pillar imparts a π -phase shift on the electric field, with respect to the adjacent vacuum space (Fig. 1b inset). As a result, electrons launched at the correct optical phase (laser–electron timing) remain phase-synchronous and experience net energy gain. In our experiment, the electron beam has a duration of 429 ± 31 fs, much longer than the laser optical cycle ($\lambda/c = 2.6$ fs, where c is the velocity of light). Electrons therefore sample all phases of the laser field, causing the laser–electron interaction to manifest as an energy modulation (broadening), with some electrons gaining energy from acceleration while some are decelerated.

Figure 1b depicts a diagram of the experiment performed at the Next Linear Collider Test Accelerator (NLCTA) facility at the SLAC National Accelerator Laboratory. The incident 60-MeV electron beam is focused by a permanent magnetic quadrupole triplet to a root mean square (r.m.s.) spot size of 24 $\mu\text{m} \times 8 \mu\text{m}$. The sample contains multiple 550- μm -long (~ 687 optical periods) DLA structures with a period of 800 nm and apertures of both 500 $\mu\text{m} \times 400$ nm and 500 $\mu\text{m} \times 800$ nm. The 400-nm gap structure provides a higher acceleration gradient

¹Department of Applied Physics, Stanford University, Stanford, California 94305, USA. ²SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, California 94025, USA. ³Department of Electrical Engineering, Stanford University, Stanford, California 94305, USA. ⁴Department of Physics and Astronomy, University of California Los Angeles, Los Angeles, California 90024, USA. ⁵Tech-X Corporation, 5621 Arapahoe Avenue, Boulder, Colorado 80303, USA.

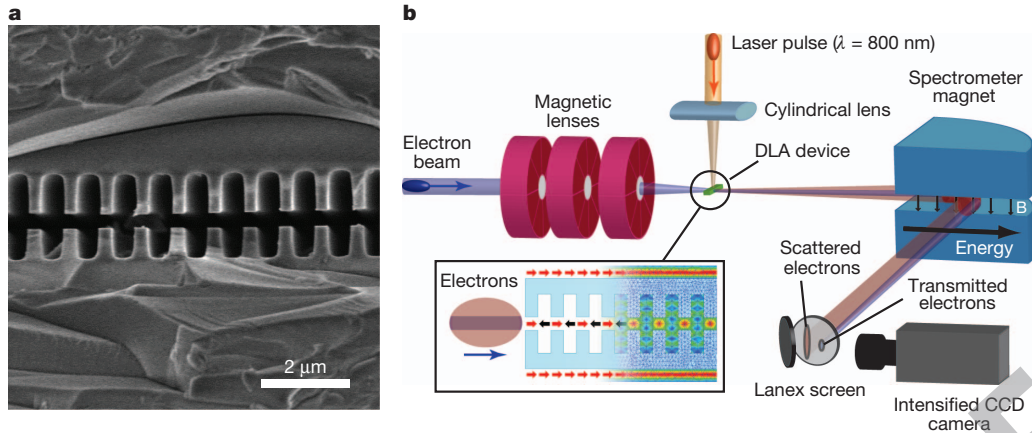


Figure 1 | DLA structure and experimental set-up. **a**, Scanning electron microscope image of the longitudinal cross-section of a DLA structure fabricated as depicted in Extended Data Fig. 1a. Scale bar, 2 μm . **b**, Experimental set-up. Inset, a diagram of the DLA structure indicating the

than an 800-nm gap structure but requires tighter tolerances on the electron beam.

The NIR pulses, 1.24 ± 0.12 ps long, from a regeneratively amplified Ti:sapphire mode-locked laser are focused to an r.m.s. spot size of $30 \mu\text{m} \times 300 \mu\text{m}$ at the interaction point. We use a motorized four-axis stage for precise alignment of the structure with the electron beam. Once aligned, the electron beam leaving the structure goes through a point-to-point focusing spectrometer magnet, which disperses the outgoing electron beam in energy onto a Kodak Lanex phosphor screen that is imaged by an intensified CCD (charge-coupled device) camera. According to particle tracking simulations, 2.2% of the 60-MeV beam is transmitted through the vacuum channel of the 400-nm gap structure (see Methods). A segment of the spectrometer screen focusing on this transmitted distribution is shown in Fig. 2a. The horizontal axis represents beam energy, and the entire image spans 240 keV. The central pixel location of the 60-MeV beam is taken as the reference point, corresponding to zero energy deviation (ΔE).

The spectrometer image in Fig. 2a is a median filtered average of a dozen shots. The least-squares fit to the distribution of electrons scattered by the fused silica substrate and the grating teeth has been removed from this image to emphasize the transmitted distribution (see Methods). A similarly averaged set of laser-on spectrometer images within 0.5 ps of the optimal timing overlap for laser pulses with energy 93 μJ per pulse is shown in Fig. 2b. The white contour in both Fig. 2a and Fig. 2b denotes the location where the spectral charge density is 4.5% of the maximum density (at the peak of the scattered distribution). In the presence of a laser field, there exists a higher charge density on either side of the original peak at $\Delta E = 0$. The white contour shows a sizable fraction of electrons with maximum energy that is ~ 60 keV higher than in the laser-off case.

The laser-induced energy modulation is readily apparent in the energy spectra (Fig. 2c). Using the fits to these spectra, a maximum energy shift of 53.1 keV is calculated from the abscissa of the half-width at half-maximum (HWHM) point in the high-energy tail. We use an analytical interaction model (see Methods and Extended Data Fig. 3) to calculate an accelerating gradient from this measurement. Figure 2c shows the input electron beam distribution used in the model (blue curve), which is a fit to the measured spectrum in the absence of a laser field (light blue crosses). The calculated energy modulation (red curve) agrees with our measurement (pink crosses), and gives a corresponding accelerating gradient of 151.2 MeV m^{-1} for this example. Particle tracking simulations (black dots; see Methods) at this gradient level give an independent confirmation of the observed modulated spectrum.

To determine the maximum gradient at a given laser power level, we measure the energy modulation as the laser pulse is temporally scanned

field polarization direction and the effective periodic phase reset, depicted as alternating red (acceleration) and black (deceleration) arrows. A snapshot of the simulated fields in the structure shows the corresponding spatial modulation in the vacuum channel. See text for details.

across the electron beam, forming a cross-correlation signal. A sample measurement at a laser pulse energy of $91.8 \pm 1.3 \mu\text{J}$ over a laser delay of 6 ps is shown Fig. 3a. The orange circles (laser-off data) show no variation correlated with laser delay, as expected, and have an r.m.s. deviation of 4.5 keV, which is taken as the noise floor level of the measurement. The blue circles (laser-on data) show the expected sech^2 distribution with a full-width at half-maximum (FWHM) of 1.89 ± 0.09 ps,

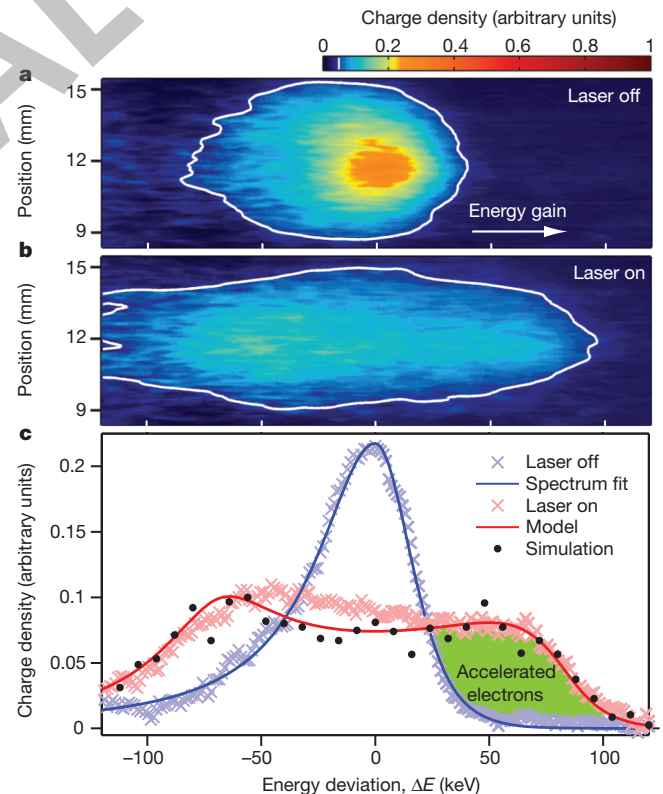


Figure 2 | Demonstration of energy modulation. **a**, Image of the transmitted electron beam on the spectrometer screen, with the laser off. **b**, As **a** but when the laser field is present. **c**, Energy spectra from **a** and **b** showing energy modulation. A fit (blue curve) to the measured laser-off spectrum (light blue crosses) is used as input for the simulations. The calculated energy modulation (red curve) and particle tracking simulations (black dots) agree with our measured spectrum (pink crosses). Images of the entire spectrometer screen are shown in Extended Data Fig. 2.

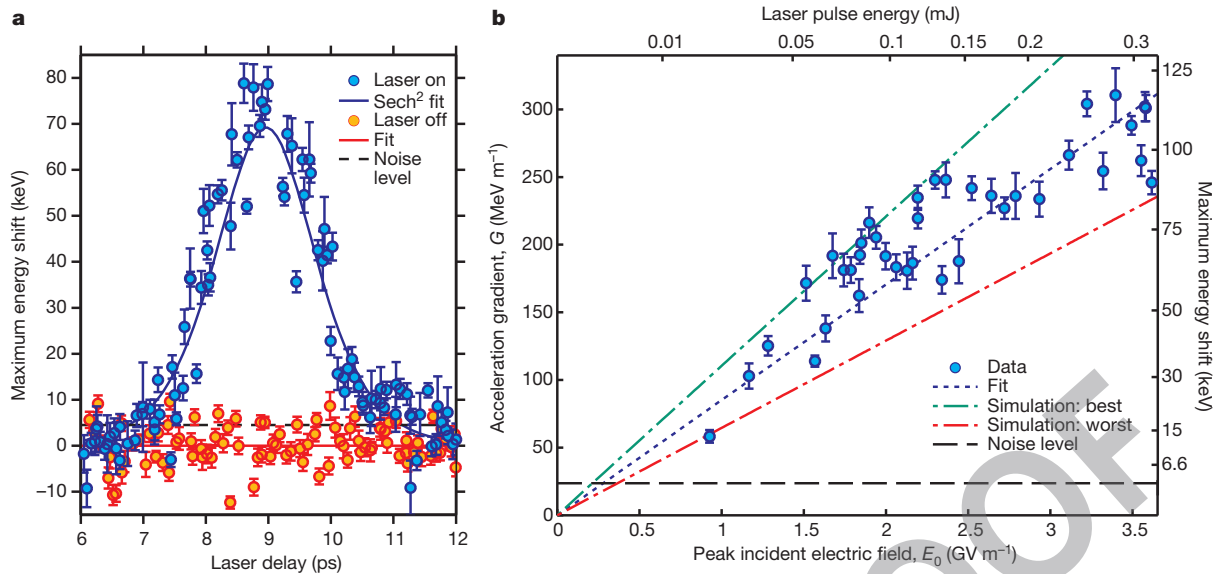


Figure 3 | Measurement of acceleration gradient in the 400-nm gap structure. **a**, Maximum energy shift as a function of laser delay give a cross-correlation measurement of the laser–electron interaction. The orange filled circles (laser-off data) show no variation with laser delay and have an r.m.s deviation of 4.5 keV which is taken as the noise level (black dashed line). The blue filled circles (laser-on data) show the expected sech^2 distribution (blue curve). **b**, Calculated acceleration gradient (blue filled circles), G , as a function

of peak incident longitudinal electric field, E_0 , showing expected linear dependence (dashed blue line). Gradients in excess of 250 MeV m^{-1} are observed within the simulated best-case (dashed green line) and worst-case (dashed red line) structure fabrication tolerance limits. Measurements on the 800-nm gap structure are shown in Extended Data Fig. 4. Error bars in **a** and **b**, 68% confidence interval.

which is in agreement with the predicted value of $2.00 \pm 0.13 \text{ ps}$ (see Methods). For each data set, the magnitude of the electron–laser interaction is determined by the amplitude of the sech^2 fit, which represents the maximum energy gain at an optimal temporal overlap. We use this figure to calculate the accelerating gradient. For the data set shown in Fig. 3a, the observed maximum energy shift is $69.2 \pm 2.3 \text{ keV}$ and the resulting gradient is $191.1 \pm 5.7 \text{ MeV m}^{-1}$, which falls within the expected range (see Supplementary Information).

The cross-correlation measurement was repeated at multiple pulse energies up to $319 \mu\text{J}$, the maximum pulse energy available. Figure 3b shows the acceleration gradient G as a function of peak incident field E_0 , along with analytical gradient calculations (see Methods) for the best case (green line) and worst case (red line) structure fabrication. We find that the measured gradient varies linearly with incident electric field as expected, and the measured values fall within the anticipated range. The highest gradient observed here is $309.8 \pm 20.7 \text{ MeV m}^{-1}$; this is twice the highest recorded values in radio-frequency (RF)-based demonstration cavities, and 6–10 times higher than typical operating gradients in RF structures. The deviation of the high-energy data points in Fig. 3b below the predicted linear trend is probably due to structural damage. Further optimization of the fabrication and materials can allow for operation at higher gradients below the laser damage fluence level. Additionally, despite week-long exposure of the structure, no electron-beam-induced charging or radiation damage was observed in the post-examination, demonstrating the robustness of dielectric structures as accelerators¹⁷.

To further verify the linear dependence of the laser–electron interaction on the electric field, the laser polarization and incidence angles were varied at a fixed laser power level. As shown in Extended Data Fig. 5, the measured gradients show good agreement with the expected dependencies and confirm that an intensity-dependent ponderomotive force is not responsible for the observed laser–electron interaction. Additionally, the fit to the data reveals that the laser–electron interaction occurs over $N = 563 \pm 104$ periods of the structure.

We have produced the first demonstration of scalable laser-driven acceleration with gradients exceeding 250 MeV m^{-1} in an ultra-compact and inexpensive DLA. The results are in agreement with both the

analytical treatment and particle tracking simulations. The pulse duration of the electron beam used in this experiment is longer than the optical cycle of the laser, which results in a modulation of the electron beam in energy. An ideal DLA-matched electron source should not only be sub-femtosecond in duration, but also have a nanometre-scale spot size, to efficiently utilize the maximum acceleration gradients in an optimized sub-wavelength aperture DLA accelerator. To create such a source, a method of bunching the electron beam at optical frequencies has been previously demonstrated¹⁸ and needle cathodes are being developed to create high-brightness low-emittance laser-triggered electron beam sources¹⁹. By bunching the beam, the fraction of captured electrons can jump from a few tens of per cent in a 5% energy window to $>80\%$ with an energy spread of $<1\%$ (ref. 20).

Integrating these novel electron sources with DLA structures would allow phase-stable net acceleration of synchronous attosecond micro-bunches over many accelerating stages. Additionally, the relatively small operating beam charge in these microstructures can be compensated by the MHz repetition rate of current laser systems to achieve the beam luminosity required for high-energy physics research²¹. Power efficiency estimates for such DLA devices are comparable with conventional RF technology, assuming that similar power efficiency (near 100%) for guided wave systems can be achieved, 40% laser wall plug efficiencies (feasible with solid state thulium fibre laser systems)²² and 40% laser to electron-beam coupling²³. Continuously tunable radiation is a natural by-product of accelerator technology, so this demonstration supports the viability of DLA for the development of compact sources of coherent, attosecond-scale X-ray pulses²⁴ with applications in medical therapy, biological and materials research, and industrial processing.

METHODS SUMMARY

The DLA structure was fabricated by bonding two fused-silica wafers, which were processed via standard optical lithography and reactive ion etching techniques¹⁵ (Extended Data Fig. 1). We used a conventional linear accelerator to generate a relativistic electron beam, with a $10 \mu\text{m}$ spot size. Owing to the dimensional mismatch between the electron beam and the sub-micrometre DLA aperture, many of the electrons experienced collisional straggling energy loss²⁵ as they traversed the fused-silica substrate (rather than the vacuum aperture). This process generated observable optical transition radiation (OTR) as a by-product, which was subsequently

used as a diagnostic to align the laser pulse to the electron beam, both spatially and temporally. The alignment of the electron beam to the aperture of the DLA structure was optimized by maximizing the population of transmitted electrons, as indicated by the electron energy spectrum (Extended Data Fig. 2). Our understanding of the observed spectrum was verified with G4Beamline²⁶ simulations of the beam scattering and radiative losses. To quantify the maximum energy gain in our structure, we developed an analytical model (Extended Data Fig. 3) and performed particle tracking simulations, both of which agreed with our measurements. Damage threshold measurements²⁷ of identical DLA structures yielded a laser damage threshold of $0.85 \pm 0.14 \text{ J cm}^{-2}$.

Online Content Any additional Methods, Extended Data display items and Source Data are available in the online version of the paper; references unique to these sections appear only in the online paper.

3 Received 28 June; accepted 16 September 2013.

Published online 27 September 2013.

- Lin, X. E. Photonic band gap fiber accelerator. *Phys. Rev. Spec. Top. Accel. Beams* **4**, 051301 (2001).
- Mizrahi, A. & Schachter, L. Optical Bragg accelerators. *Phys. Rev. E* **70**, 016505 (2004).
- Cowan, B. M. Three-dimensional dielectric photonic crystal structures for laser-driven acceleration. *Phys. Rev. Spec. Top. Accel. Beams* **11**, 011301 (2008).
- Plettner, T., Lu, P. P. & Byer, R. L. Proposed few-optical cycle laser-driven particle accelerator structure. *Phys. Rev. Spec. Top. Accel. Beams* **9**, 111301 (2006).
- Ginzton, E. L., Hansen, W. W. & Kennedy, W. R. A linear electron accelerator. *Rev. Sci. Instrum.* **19**, 89–108 (1948).
- Plettner, T. *et al.* Visible-laser acceleration of relativistic electrons in a semi-infinite vacuum. *Phys. Rev. Lett.* **95**, 134801 (2005).
- Shimoda, K. Proposal for an electron accelerator using an optical maser. *Appl. Opt.* **1**, 33–35 (1962).
- Piestrup, M. A., Rothbart, G. B., Fleming, R. N. & Pantell, R. H. Momentum modulation of a free electron beam with a laser. *J. Appl. Phys.* **46**, 132–137 (1975).
- Palmer, R. B. Interaction of relativistic particles and free electromagnetic waves in the presence of a static helical magnet. *J. Appl. Phys.* **43**, 3014–3023 (1972).
- van Steenbergen, A., Gallardo, J., Sandweiss, J. & Fang, J.-M. Observation of energy gain at the BNL inverse free-electron-laser accelerator. *Phys. Rev. Lett.* **77**, 2690–2693 (1996).
- Marshall, T. C. *Free-Electron Lasers* 24–26 (McMillan, 1985).
- Takeda, Y. & Matsui, I. Laser linac with grating. *Nucl. Instrum. Methods* **62**, 306–310 (1968).
- Mizuno, K., Pae, J., Nozokido, T. & Furuya, K. Experimental evidence of the inverse Smith-Purcell effect. *Nature* **328**, 45–47 (1987).
- Breuer, J. & Hommelhoff, P. Laser-based acceleration of non-relativistic electrons at a dielectric structure. *Phys. Rev. Lett.* (in the press).
- Peralta, E. A. *et al.* Design, fabrication, and testing of a fused-silica dual-grating structure for direct laser acceleration of electrons. *AIP Conf. Proc.* **1507**, 169–177 (2012).
- Plettner, T., Byer, R. L. & Montazeri, B. Electromagnetic forces in the vacuum region of laser-driven layered grating structures. *J. Mod. Opt.* **58**, 1518–1528 (2011).
- Thompson, M. C. *et al.* Breakdown limits on gigavolt-per-meter electron-beam-driven wakefields in dielectric structures. *Phys. Rev. Lett.* **100**, 214801 (2008).
- Sears, C. M. S. *et al.* Production and characterization of attosecond electron bunch trains. *Phys. Rev. Spec. Top. Accel. Beams* **11**, 061301 (2008).
- Ganter, R. *et al.* Laser-photofield emission from needle cathodes for low-emittance electron beams. *Phys. Rev. Lett.* **100**, 064801 (2008).
- Duris, J. P., Musumeci, P. & Li, R. K. Inverse free electron laser accelerator for advanced light sources. *Phys. Rev. Spec. Top. Accel. Beams* **15**, 061301 (2012).
- Colby, E. R., England, R. J. & Noble, R. J. A laser-driven linear collider: sample machine parameters and configuration. *Proc. 24th Particle Accelerator Conf. C110328* 262–264 (2011).
- Moulton, P. F. *et al.* Tm-doped fiber lasers: fundamentals and power scaling. *IEEE J. Sel. Top. Quant.* **15**, 85–92 (2009).
- Siemman, R. Energy efficiency of laser driven, structure based accelerators. *Phys. Rev. Spec. Top. Accel. Beams* **7**, 061303 (2004).
- Plettner, T. & Byer, R. L. Microstructure-based laser-driven free-electron laser. *Nucl. Instrum. Methods A* **593**, 63–66 (2008).
- Warner, C. & Rohrlich, F. Energy loss and straggling of electrons. *Phys. Rev.* **93**, 406–407 (1954).
- Roberts, T. J. & Kaplan, D. M. G4beamline simulation program for matter-dominated beamlines. *Proc. 22nd Particle Accelerator Conf. C070625* 3468–3470 (2007).
- Soong, K., Byer, R. L., Colby, E., England, R. J. & Peralta, E. A. Laser damage threshold measurements of optical materials for direct laser accelerators. *AIP Conf. Proc.* **1507**, 511–515 (2012).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank R. Noble, J. Spencer, O. Solgaard and J. Harris for discussions, J. Nelson, D. McCormick and K. Jobe for technical assistance at SLAC, and M. Tang, M. Mansourpour, N. Latta, M. Stevens, J. Conway and U. Thumser for technical assistance at the Stanford Nanofabrication Facility (SNF). This work was supported by the US DoE (grant no. DE-FG03-92ER40693) and DARPA (grant no. N66001-11-1-4199). Device fabrication took place at SNF, which is supported by the NSF under grant ECS-9731293.

Author Contributions E.A.P., K.S., R.J.E., E.R.C. and R.L.B. designed the experiment. E.A.P., K.S., R.J.E., E.R.C., Z.W. and B.M. built the experiment, and with the help of C.McG. and J.McN., carried out the experiment. E.A.P. designed and fabricated the structure. E.A.P. and K.S. wrote the data analysis software and analysed the results, and E.B.S. contributed to that software. K.S. developed the analytical model. K.S., B.M., J.McN, B.C. and B.S. performed supporting simulations. G.T, J.McN, E.B.S. and K.J.L. provided feedback to improve the experiment. K.S. and K.J.L. designed and constructed the vacuum damage threshold measurement set-up, and K.S. carried out such measurements. D.W. performed hardware upgrades and installation. E.A.P. wrote the manuscript with contributions from K.S. and R.J.E. and revisions by all.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to R.L.B. (rlbyer@stanford.edu).

METHODS

Grating structure fabrication. Fabrication of the dual layer grating structures was carried out at the Stanford Nanofabrication Facility (SNF). The gratings are patterned on fused silica wafers via optical lithography using an ASML PAS 5500 i-line stepper and magnetically enhanced reactive ion etching on an Applied Materials P5000 etcher. This is shown schematically in steps 1–3 of Extended Data Fig. 1a.

Four grating structures are included in every sample to facilitate a switch during the experiment, two with a 400-nm gap, and two with an 800-nm gap. Once patterned and etched, alignment channels of decreasing sizes ($250\ \mu\text{m} \times 250\ \mu\text{m} \rightarrow 80\ \mu\text{m} \times 80\ \mu\text{m} \rightarrow 80\ \mu\text{m} \times 20\ \mu\text{m}$) are cut into the wafers to facilitate the alignment of the electron beam with the structure. Two wafers are then bonded directly to each other via fusion bonding and diced into individual samples (steps 4–5 of Extended Data Fig. 1a). The alignment tolerance of the bonding process is $\sim 3\ \mu\text{m}$ so the grating-to-grating longitudinal alignment is not ensured. However, a perfectly aligned structure is not required to generate substantial acceleration gradients¹⁵. As a final step, a section of the sample is metal-coated to serve as a reflector for the NIR laser, and a wedge is cut at 45° to be able to image the infrared laser beam and optical transition radiation (OTR) simultaneously. These two features, along with the alignment channels and the four grating structures, are shown in the schematic in Extended Data Fig. 1c.

Electron beam. Testing of the dielectric accelerator structures is conducted using pre-accelerated electron bunches at the Next Linear Collider Test Accelerator (NLCTA) facility at SLAC National Accelerator Laboratory. The electron bunches are produced by an RF photoinjector and accelerated up to 60 MeV in a travelling wave RF cavity at a repetition rate of 10 Hz. Energy collimation is used to create a stable monoenergetic ‘slice’ of the electron bunch with a FWHM energy spread of 24 keV or 0.04%, an r.m.s. length of $129 \pm 9\ \mu\text{m}$ ($0.43 \pm 0.03\ \text{ps}$), and emittances of $16\ \text{mm mrad} \times 7\ \text{mm mrad}$ in X and Y , respectively. Of the initial 5 pC of charge, 10% is transmitted through the energy collimators, giving an estimated charge of 0.5 pC incident on the grating structure. Based on particle scattering simulations, 2.2% of these particles are transmitted through the 400-nm gap structure’s vacuum channel. Consequently, the population of electrons that pass through the accelerating channel of the 400-nm gap device contains an estimated 11 fC of charge per electron bunch. A sample spectrometer image for the case of the electron beam travelling entirely through vacuum (that is, in the absence of the structure) is shown in Extended Data Fig. 2a. The horizontal axis represents beam energy and the resolution limit of the spectrometer is 1.2 keV per image pixel so the entire image spans 800 keV. The central pixel location of the narrow energy spread 60-MeV beam is taken as the reference point, corresponding to zero energy deviation (ΔE).

Achieving electron transmission. Transmission of electrons through the sub-micrometre grating aperture is achieved with the help of alignment apertures of decreasing size that run parallel to the grating vacuum channels (see Extended Data Fig. 1c). Once transmission on the large channel is observed, the structure’s optimal pitch and yaw are determined by maximizing the transmitted distribution. This is repeated for the smaller apertures until we see transmission through the grating aperture. Once this is achieved, the laser beam incidence angle is adjusted to ensure perpendicularity.

Because the NLCTA beam is significantly larger than the structure aperture, only a small fraction of electrons travels through the vacuum channel when the DLA structure is in place. Extended Data Fig. 2b shows the corresponding spectrometer image for this case, containing two distinct distributions. The prominent broad distribution, centred at $\Delta E = -340\ \text{keV}$, represents the (majority) electron population that travels primarily through the fused-silica substrate. The energy of these electrons is reduced due to collisional straggling loss in the material²⁵. The smaller, more compact distribution at $\Delta E = 0\ \text{keV}$ corresponds to the population of electrons that traverse the grating structure primarily through the vacuum channel. Because the channel has a 1,280:1 aspect ratio (430 nm tall, and 550 μm long), these electrons are collimated with an acceptance angle of 0.04° . This distribution therefore has a reduced vertical spread but retains its original energy spread. Henceforth, we refer to the higher-energy peak as the transmitted population and the lower-energy peak as the scattered population.

The formation of these two distributions is confirmed by simulations of the scattering and radiative losses of the electron beam as it traverses the grating structure. Extended Data Fig. 2c shows a projection of the spectrometer image of Extended Data Fig. 2b onto the energy axis which gives the resulting electron energy spectrum. Also shown is the calculated spectrum from particle scattering simulations, which has peak locations and relative amplitudes in agreement with the observed experimental data. Because the accelerating fields are strongest inside the vacuum channel, we focus on the transmitted distribution when studying the effects of the laser–electron interaction. In our analysis, the resulting spectra are fitted to a curve composed of two distributions (scattered and transmitted), using

the method of least-squares. A sample fit, shown (in orange) in Extended Data Fig. 2c, matches the measured spectra closely.

Particle scattering simulations. We model the scattering and radiative losses of the electron beam as it traverses the grating structure using the Geant4 code G4Beamline²⁶ to calculate the expected electron energy spectrum. In this model, an electron beam with a transverse spot size of $24\ \mu\text{m} \times 8\ \mu\text{m}$ and a transverse vertical emittance of 3 mm mrad is sent through an approximate model of the grating structure with a 400-nm gap. The model structure is made up of five layers. A central, 400-nm tall, vacuum layer surrounded by two 660-nm-thick grating layers in the first 550 μm of the longitudinal dimension. Two 1.1-mm-long, 500- μm -thick layers of fused silica surround the grating layers. Extended Data Fig. 2c shows the results of this simulation in agreement with the experimental data, giving the correct spectrum peak separation of 340 keV and relative peak amplitude.

The simulations give a calculated percentage of 2.2% of the total number of electrons in the transmitted distribution. By comparison, simulations of the 800-nm gap structure yield 5.8% of the electrons in the transmitted distribution.

Energy modulation simulations. The fields in the structure vacuum channel are calculated using the finite-element frequency-domain code HFSS and imported into a MATLAB-based code that computes the resulting relativistic Lorentz force on a test particle as it traverses the structure. In our simulations, we approximate the NLCTA electron beam using 2,500 test particles, whose position and momentum match those of the experiment. Owing to the long computation time required to perform such particle tracking simulations, we elect to model and scale a 100-period structure rather than model the entire 650-period structure. Based on convergence tests, we find that this approach produces accurate results, which is expected from the linearity of the acceleration effect, at a fraction of the computation cost. The final energy distribution of the 2,500-particle ensemble is then calculated from the results of each individual particle trace.

Analytical treatment of energy modulation. The interaction between a NIR laser pulse and a picosecond-scale electron bunch in a grating accelerator structure can be described analytically by partitioning the long electron bunch into a series of short finite slices of length τ , as shown in Extended Data Fig. 3a, where each slice satisfies the relation $\tau \ll \lambda/c$. In the presence of the sinusoidally varying electric fields of the structure, each slice of the electron bunch will experience a net energy shift with a negligible effect on the energy profile. For example, a slice of an electron bunch with an asymmetric energy profile typical of the NLCTA will experience a net energy gain of 50 keV if injected at the optimal phase in a millimetre-long structure capable of $50\ \text{MeV m}^{-1}$ gradients (see green slice in Extended Data Fig. 3b). The net effect of the structure is calculated by superimposing the effect for each individual slice. For a long pulse, we find that the predicted effect on the electron energy spectrum is a transition from a single peaked distribution to a double peaked profile of lower amplitude, as shown in Extended Data Fig. 3c.

Energy spectrum fits. The projected energy spectra are fitted to curves containing the two distinct, scattered and transmitted, electron distributions. Several different profiles including asymmetric Gaussian, Lorentzian, sech^2 and Landau distributions, along with combinations of these were initially employed. These were compared against the measured single-peak data (that is, electron beam incident on the fused-silica sample but away from the grating vacuum channel) to determine a profile that minimized the χ^2 value of the fits. The scattered distribution was ultimately fitted to a half Lorentzian on its low-energy side, and a half sech^2 on its high-energy side. The transmitted distribution was fitted to a double-peak version of the scattered distribution, with the separation of the two peaks as a free parameter and mirrored about its centre. The sum of the two distributions fit the observed spectra very well, as indicated in Extended Data Fig. 2c.

On the basis of comparisons with the particle scattering simulations and the analytical treatment, we find that a partially-straggled population exists, composed of electrons which are scattered (partially) by the grating teeth. This is accurately accounted for by the low-energy side Lorentzian function used to fit the transmitted distribution. To isolate the fully transmitted electrons (as shown in Fig. 2c), we subtract both the straggled distribution and the partially straggled distribution. The difference between the abscissa of the HWHM point in the high-energy tail for a laser-on spectrum and a laser-off spectrum gives a measurement of the laser-induced maximum energy shift.

Spatial overlap. We use a miniature aluminium pellicle mirror and a long working distance Cassegrain telescope to image the back of the grating sample with a resolution limit of $8\ \mu\text{m}$ per image pixel. To perform the spatial alignment, we simultaneously determine the location of the OTR origin on the sample and the laser beam scattering off the fabricated 45° wedge shown in Extended Data Fig. 1c. The laser beam is then steered until the observed scattering lines up with the OTR location to within $\sim 10\ \mu\text{m}$.

Temporal overlap. A fast photodiode (25-GHz bandwidth) is used to detect the OTR signal the electron beam generates at an aluminium pellicle located downstream of the sample. The transmitted NIR pulse is also measured simultaneously,

on the same detector and oscilloscope. We observe the measured arrival time of the two signals and adjust a four-pass optical delay line to achieve sub-nanosecond overlap. A motorized delay line provides gross overlap to within the 50-ps resolution limit of the oscilloscope, and a voice coil actuated retroreflector is scanned during the experiment to achieve picosecond-scale timing.

Cross-correlation measurement. A typical cross-correlation signal in our experiment consist of 300 independent electron-laser interaction events measured at 150 unique temporal overlap positions. The resulting cross-correlation signal follows a hyperbolic secant squared (sech^2) pattern as expected based on the laser pulse profile indicated by autocorrelation measurements.

Typical cross-correlation measurements co-propagate two signals through an interacting medium, while varying the temporal overlap. In this case, the cross-correlation signal is solely dependent on the temporal dimension. However, in our cross-correlation measurement, we use two signals incident at perpendicular angles, and as a result, the cross-correlation signal is a function of both the temporal and spatial dimensions of our electron bunch and our laser beam. The expected outcome of this interaction can be calculated by simulating the perpendicular crossing of an electron bunch with a laser beam, with the temporal and spatial dimensions set to match the experimental values. From this calculation, we predict a FWHM of 2.00 ± 0.13 ps for the sech^2 fit to the cross-correlation signal. This is in strong agreement with our experimental findings, where a cross-correlation FWHM equal to 1.89 ± 0.09 ps was observed.

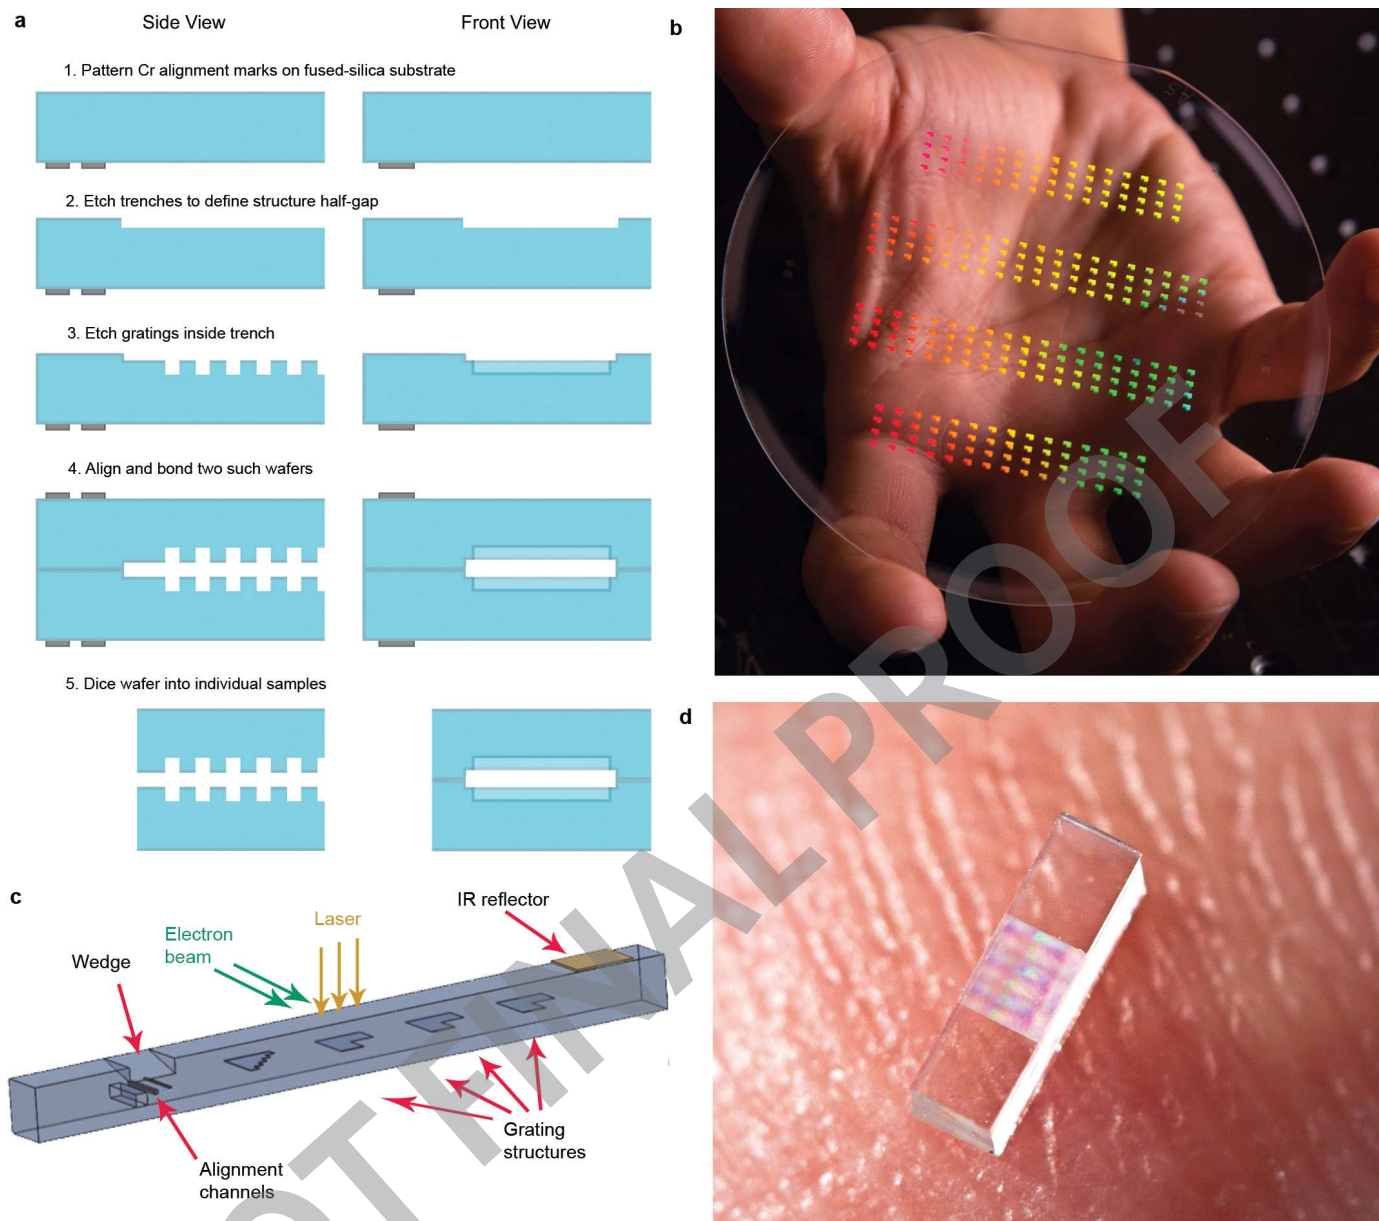
Damage threshold measurement. Damage threshold measurements of DLA structures from the same wafer as the one used in the acceleration experiment were performed using an experimental set-up similar to the one described in ref. 27.

An array of accelerator structures was placed in a vacuum chamber at a pressure of 10^{-5} torr. The NIR pulses from our Ti:sapphire laser system was then focused to a r.m.s. spot size of $27 \mu\text{m} \times 53 \mu\text{m}$ and incident on the structure, in the same manner as in the acceleration experiment. The laser fluence was then gradually increased from an initial value of $2.00 \mu\text{J}$ in increments of $1.85 \mu\text{J}$, while the illuminated site was monitored with a camera. At each energy level, approximately 5,000 laser pulses were applied to the structure. Damage was determined to occur when a visible change in the strength and consistency of the scattered infrared light at the illumination site was observed. A total of 24 independent damage threshold measurements were taken, and each damage site was afterwards inspected under a high powered optical microscope to verify the occurrence of laser damage.

The results of these measurements indicate a laser damage threshold of $0.85 \pm 0.14 \text{ J cm}^{-2}$ for our grating accelerator structures. From post-examination of the structure it is clear that this damage threshold value was exceeded during the course of the experiment. However, we continued to see acceleration up to the maximum available pulse energy. On the basis of post-analysis of a laser damage site, we deduced that our laser r.m.s. spot size during the experiment was actually $76 \mu\text{m} \times 294 \mu\text{m}$, yielding a maximum fluence of 0.85 J cm^{-2} on the gratings.


Error analysis. The corresponding error bar for each data point represents a 68% confidence interval for the plotted fit parameter, as calculated by the Jacobian of the weighted nonlinear least-squares regression. The noise floor in the plots is defined by the laser-off data for each set of measurements. Specifically, we define the gradient of the laser-off data as zero gradient, and used the positive tail of the 68% confidence interval as the noise floor.

NOT FINAL PROOF

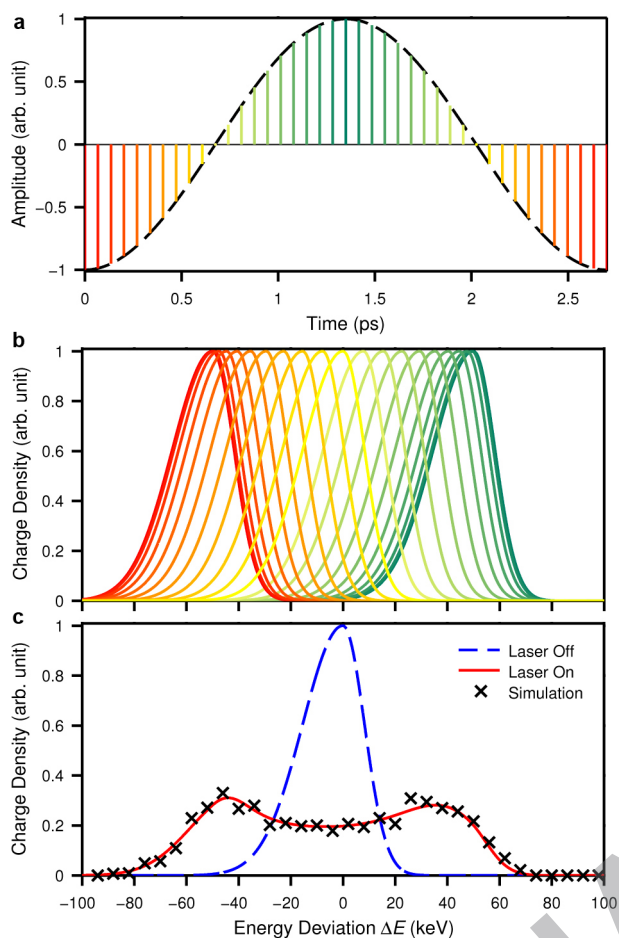


Extended Data Figure 1 | DLA sample preparation. **a**, Diagram of the structure fabrication process. In side view, the electron beam traverses the structure from left to right. In front view, the beam goes into the page. Laser is incident from above. See Methods for details. **b**, Picture of a completed wafer with hundreds of DLA structures. **c**, Diagram of a finished sample ready for

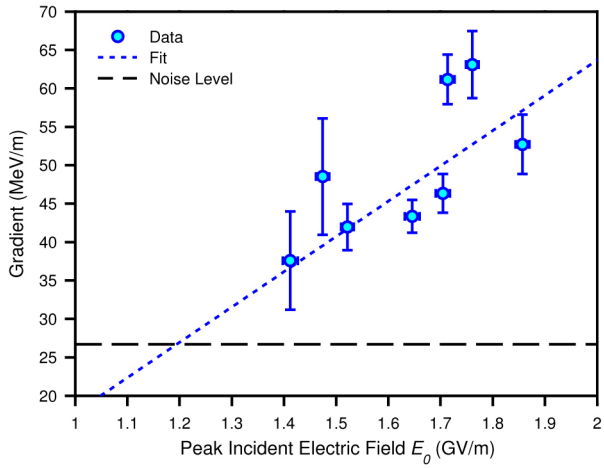
beam tests including four DLA structures, alignment channels, a wedge for spatial alignment of the laser to the electron beam, and a metal coating for perpendicular alignment of the laser. **d**, Picture of a single DLA structure on a fingertip.



Extended Data Figure 2 | Reference spectrometer screen images and spectrum. **a**, Spectrometer screen image showing the 60-MeV beam with no DLA structure in place. **b**, Spectrometer screen image of the beam after traversing the grating structure. **c**, Projection of **b** onto the energy axis yields the energy spectrum (black crosses) in agreement with particle scattering simulations (blue dots). The corresponding least squares spectrum fit (orange curve) is also shown.

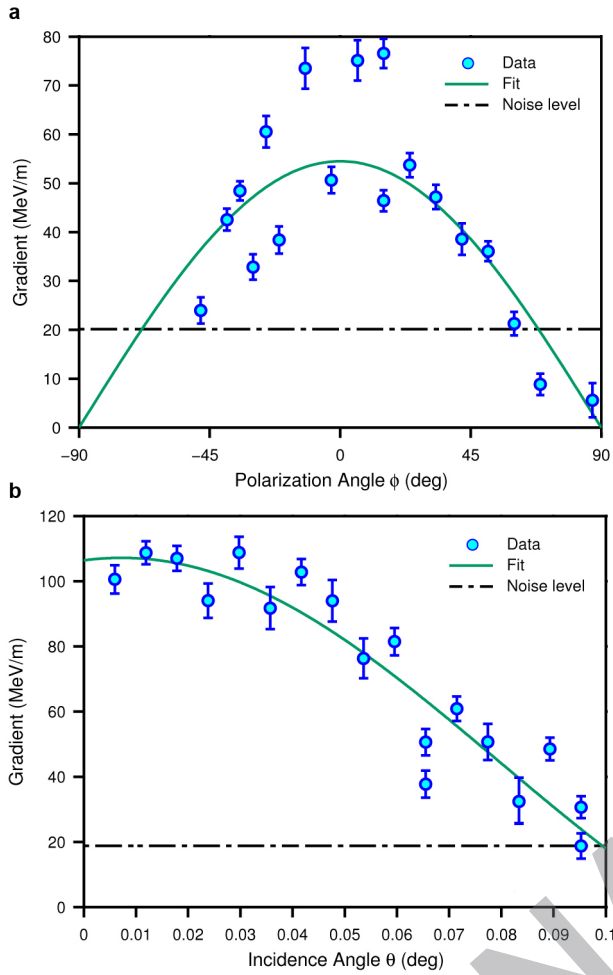


Extended Data Figure 3 | Analytical treatment of laser-driven electron energy modulation. **a**, Picosecond-scale electron beam partitioned into a series of finite slices over one optical cycle of the laser (that is, the green slice is at optimal phase for acceleration). **b**, Each slice samples a different phase of the laser pulse and therefore experiences a corresponding net energy shift with a negligible effect on the energy profile (that is, the green slice experiences net energy gain). **c**, When all contributions are superimposed, the initial single distribution (dashed blue line) becomes a double-peaked profile (red line), in agreement with particle tracking simulations (black crosses).



Extended Data Figure 4 | Gradient measurement on 800-nm gap structure. Calculated gradient (blue filled circles) G as a function of longitudinal electric field E_0 showing expected linear dependence (dashed blue line) and reduced strength when compared to the 400-nm gap structure, as expected for a larger gap. The dashed black line is the measurement noise level. Error bars, 68% confidence interval.

NOT FINAL PROOF



Extended Data Figure 5 | Experimental verification of direct laser acceleration. **a**, Gradient in the 800-nm gap structure for an input pulse energy of $105.2 \pm 3.6 \mu\text{J}$. As the laser polarization is rotated away from the direction of electron propagation by an angle ϕ , the acceleration gradient varies as $G \propto \cos\phi$. Data for $\phi < 0$ were taken last, and beam quality had degraded. **b**, Gradient in the 400-nm gap structure at an input pulse energy of $29.3 \pm 0.4 \mu\text{J}$, averaging the observed modulation over many shots taken at an optimal timing overlap. As the laser incidence angle deviates from perpendicular by an angle θ , the observed gradient decreases according to the expected relationship, equation (2) in Supplementary Information. Data are shown as blue filled circles with the corresponding least squares fit shown as green lines. The dashed black line is the measurement noise level. Error bars, 68% confidence interval.