



Revealing the three-dimensional structure of microbunched plasma-wakefield-accelerated electron beams

Received: 24 January 2024

Accepted: 6 June 2024

Published online: 15 July 2024

Check for updates

Maxwell LaBerge ^{1,2}, Brant Bowers ¹, Yen-Yu Chang ²,
Jurjen Couperus Cabadağ ², Alexander Debus ², Andrea Hannasch¹,
Richard Pausch ², Susanne Schöbel ^{2,3}, Jessica Tiebel ^{2,3}, Patrick Ufer ^{2,3},
Anna Willmann^{2,3}, Omid Zarini^{2,3}, Rafal Zgadaj¹, Alex H. Lumpkin⁴,
Ulrich Schramm ^{2,3}, Arie Irman ² & M. C. Downer ¹✉

Plasma wakefield accelerators use tabletop equipment to produce relativistic femtosecond electron bunches. Optical and X-ray diagnostics have established that their charge concentrates within a micrometre-sized volume, but its sub-micrometre internal distribution, which critically influences gain in free-electron lasers or particle yield in colliders, has proven elusive to characterize. Here, by simultaneously imaging different wavelengths of coherent optical transition radiation that a laser-wakefield-accelerated electron bunch generates when exiting a metal foil, we reveal the structure of the coherently radiating component of bunch charge. The key features of the images are shown to uniquely correlate with how plasma electrons injected into the wake: by a plasma-density discontinuity, by ionizing high-Z gas-target dopants or by uncontrolled laser-plasma dynamics. With additional input from the electron spectra, spatially averaged coherent optical transition radiation spectra and particle-in-cell simulations, we reconstruct coherent three-dimensional charge structures. The results demonstrate an essential metrology for next-generation compact X-ray free-electron lasers driven by plasma-based accelerators.

The production of intense coherent radiation by microbunched relativistic electron (e^-) beams underlies today's X-ray free-electron lasers (XFELs)¹, which are revolutionizing our understanding of biomolecules, materials and living cells². For an e^- bunch to radiate coherently, some of its electrons must possess a space-time-correlated structure. Such microbunching emerges from noise when an e^- beam interacts with a magnetic undulator, inducing the beam to radiate partially coherent self-amplified spontaneous emission³. An e^- beam can also pre-bunch

from interacting with electromagnetic or material structures before the undulator. Such seed microbunching can profoundly influence the FEL gain dynamics and the intensity, coherence and stability of FEL radiation⁴. Hence, non-intercepting diagnostics that track the evolution of microbunching from the cathode through the undulator are critical to develop and operate accelerator-based coherent light sources. Indeed, single-wavelength measurements of optical transition radiation that e^- beams generate on transiting thin metal films⁵,

¹Department of Physics, University of Texas at Austin, Austin, TX, USA. ²Institute for Radiation Physics, Helmholtz-Zentrum Dresden-Rossendorf, Dresden, Germany. ³TUD Dresden University of Technology, Dresden, Germany. ⁴Accelerator Directorate, Fermi National Accelerator Laboratory, Batavia, IL, USA.

✉e-mail: downer@physics.utexas.edu

particularly its strong coherent component (coherent optical transition radiation (COTR))⁶, were pivotal to understand how microbunching evolved in early FELs^{7–9} and thus to the emergence of today's XFELs.

Recent observations of gain in FELs driven by e^- beams from centimetre-scale laser-driven^{10,11} and beam-driven^{12,13} wakefield accelerators (WFAs) now portend a new generation of tabletop XFELs^{14,15}. However, WFA-driven bunches differ markedly from their radio-frequency-accelerated counterparts: they are typically shorter in duration (down to a few femtoseconds), correspondingly higher in peak current (tens of kiloamperes) and emerge from the accelerator with a smaller transverse size ($\sim 1 \mu\text{m}$), properties that can promote the FEL gain. On the other hand, they typically possess larger relative energy spread ($\Delta E_e/E_e > 10^{-2}$), larger divergence (few milliradians) and higher shot-to-shot beam-pointing fluctuations ($> 0.1 \text{ mrad r.m.s.}$) (ref. 16), factors that inhibit the FEL gain. In addition, the complex laser–plasma or beam–plasma interaction that creates plasma accelerator structures can inject plasma electrons into them that pre-bunch on micrometre^{17,18} to nanometre¹⁹ scales, often with intertwined longitudinal and transverse coherence. These unique properties of WFA e^- beams underlie proposals for new types of coherent X-ray source driven by chirped^{20,21}, attosecond-compressed²² and nano-bunched¹⁹ WFA beams. These same properties demand new ways of diagnosing WFA e^- -beam microbunching²³, which will be as essential for developing WFA-driven coherent light sources as previous diagnostics^{7–9} were for developing today's radio-frequency-accelerator-driven FELs.

Here we report a minimally intercepting, single-shot diagnostic that unfolds the three-dimensional (3D) structure of the microbunched sub-population of laser-driven wakefield accelerator (LWFA) e^- bunches from the wide-bandwidth COTR that they uniquely generate. Previous work determined one-dimensional (longitudinal) profiles of such e^- bunches^{24–26} from transversely averaged, multioctave COTR spectra²⁷ analysed with an iterative reconstruction algorithm that assumed a smooth transverse structure²⁸. Similar analysis supplements some measurements presented here. However, the critical new element of the present work is the imaging of the COTR at near-infrared to ultraviolet wavelengths—that is, wavelengths close to the dimensions of the e^- bunch itself—from the COTR foil to nine high-resolution array detectors, each filtered to detect a specific wavelength λ . In this λ range, the COTR changes from nearly fully coherent at longer λ to nearly fully incoherent at shorter λ and captures the most essential information about the beam's transverse microbunched structure. Combining λ -resolved COTR images with longitudinal profiles obtained from, for example, multioctave COTR spectra, electron spectra or particle-in-cell (PIC) simulations then elucidates the beam's 3D microbunched structure. For simply structured beams, a full 3D reconstruction is possible with high confidence. For more complex beams, the contributing 3D structure can still be estimated with help from PIC simulations. In either case, a combined transverse–longitudinal analysis is essential because transverse (\perp) and longitudinal (\parallel) form factors²⁹ contribute either as a separable product $F_{\perp}(\mathbf{k}_{\perp})F_{\parallel}(k_z)$ or a correlated function $F(\mathbf{k})$ to the total beam microbunching, where $\mathbf{k} = \mathbf{k}_{\perp} + k_{\parallel}\hat{z}$ is a wavevector of magnitude $k = \|\mathbf{k}\| = 2\pi/\lambda$. Thus, oversimplified assumptions about one lead to errors in the other.

Our main result is the observation of a correlation between the size, shape, intensity and λ dependence of a multispectral COTR image set on one hand and on the other, the method of injecting plasma electrons into the LWFA. Injection by propagating the laser through a preformed plasma-density down-ramp³⁰ in a laser-ionized helium gas jet led to annular-shaped, near- λ -independent COTR images, and to simply shaped, micrometre-scale reconstructed e^- bunches. Injection by the delayed ionization of the inner-shell electrons of nitrogen dopants in the helium plasma (self-truncated ionization injection (STII)³¹) led to larger, λ -dependent COTR images with a more complex, but stable, structure. Aided by PIC simulations, we identified these

COTR features with density modulations near the beginning and end of the $\sim 5\text{-}\mu\text{m}$ -long e^- bunch. Finally, injection via uncontrolled laser–plasma dynamics in a pure, featureless He plasma (self-injection) led to the largest, most complex and least stable COTR images. Such empirical correlations lend themselves to machine learning and feedback control of LWFAs³², even when the coherently radiating e^- -beam profile cannot be uniquely reconstructed. Although the current results were obtained from a COTR foil just outside the accelerator exit, the same method can track microbunching at downstream locations. Similar to other coherent imaging techniques such as ptychography³³ and coherent diffraction imaging³⁴, multispectral COTR imaging shares the problem of iteratively reconstructing a micro-/nanostructured object of interest from the intensity measurements of coherent interference patterns. Its distinguishing features are the relativistic speed of the object of interest, which generates its own coherent radiation in a broadband ultrashort burst, and the λ multiplexing of the resulting interference pattern to form the core dataset for reconstruction.

Results

Figure 1 (top) depicts the experimental setup. Laser pulses of 67 TW ($\sim 2\text{J}$, 30 fs) peak power and $\lambda = 800 \text{ nm}$ from the Dresden laser acceleration source (DRACO) laser at Helmholtz-Zentrum Dresden-Rossendorf (HZDR) (Methods) were focused into a gas jet, where they drove non-linear plasma wakes that accelerated electrons injected from the surrounding plasma to energies as high as 300 MeV (ref. 35). A tilted 65 μm aluminium foil 2 mm after the jet blocked the transmitted laser pulses while only minimally perturbing the e^- beams^{36,37}. The e^- beams then passed through a 25 μm Kapton foil 1 mm further downstream, inducing radial currents in its 0.3 μm aluminized back surface that emitted forward COTR as the beams exited the foil. A polished 100 μm thick silicon wafer reflected the COTR towards a microscope objective, which imaged it from the COTR foil to nine bandpass-filtered charge-coupled-device cameras (a camera filtered at 193 nm, which detected no signal, is not shown in Fig. 1). The e^- beam passed through the wafer to an electron spectrometer (Fig. 1, top right). No COTR or electrons were observed with the gas jet turned off. Methods provides further details.

COTR multispectral imaging phenomenology

Experimental COTR images. The three-row bottom panel in Fig. 1 shows the representative eight-panel multispectral COTR image sets acquired in one shot for each type of injection. Down-ramp injection yielded the simplest, smallest and most consistent COTR patterns: single annuli with deep central minima and similar size for all the measured wavelengths. STII bunches yielded larger, more complex images. Self-injected bunches yielded the largest, most complex images, varying from λ to λ and shot to shot without a discernible pattern. Figure 2a plots the effective radii $r_{\text{eff}} = \sqrt{A/\pi}$ of COTR images as a function of λ for each type of injection. Here A is an area, illustrated by the dashed yellow lines in the 400 nm column in Fig. 1, containing all of the pixels with signals over 1/4 of the maximum along with any enclosed dark minima. Down-ramp injection (blue data points) yielded not only the smallest images ($r_{\text{eff}} = 9.0 \pm 2.0 \mu\text{m}$) but also the smallest variation amongst the wavelengths and from shot to shot (error bars), whereas self-injection (green) led to the corresponding largest values.

The ratio $W(\lambda)/W_1(\lambda)$ of the total energy $W(\lambda)$ deposited into each image, determined from the integrated camera counts corrected for detector quantum efficiency and optical losses, to the energy $W_1(\lambda)$ that one electron deposits, determined from the optical transition radiation theory^{5,6}, is a sensitive, quantitative indicator of a microbunched (that is, coherently radiating) e^- -beam substructure. This is because the coherent part of $W(\lambda)/W_1(\lambda) = N_i + N_c^2 \|\mathbf{F}(\mathbf{k})\|^2$ (ref. 9) is proportional to $N_c^2 \|\mathbf{F}(\mathbf{k})\|^2$, where N_c is the number of microbunched electrons contributing to COTR with $0 < \|\mathbf{F}(\mathbf{k})\|^2 \leq 1$ at each λ , whereas the incoherent optical transition radiation energy scales only linearly with the number N_i of non-microbunched electrons. Here, since

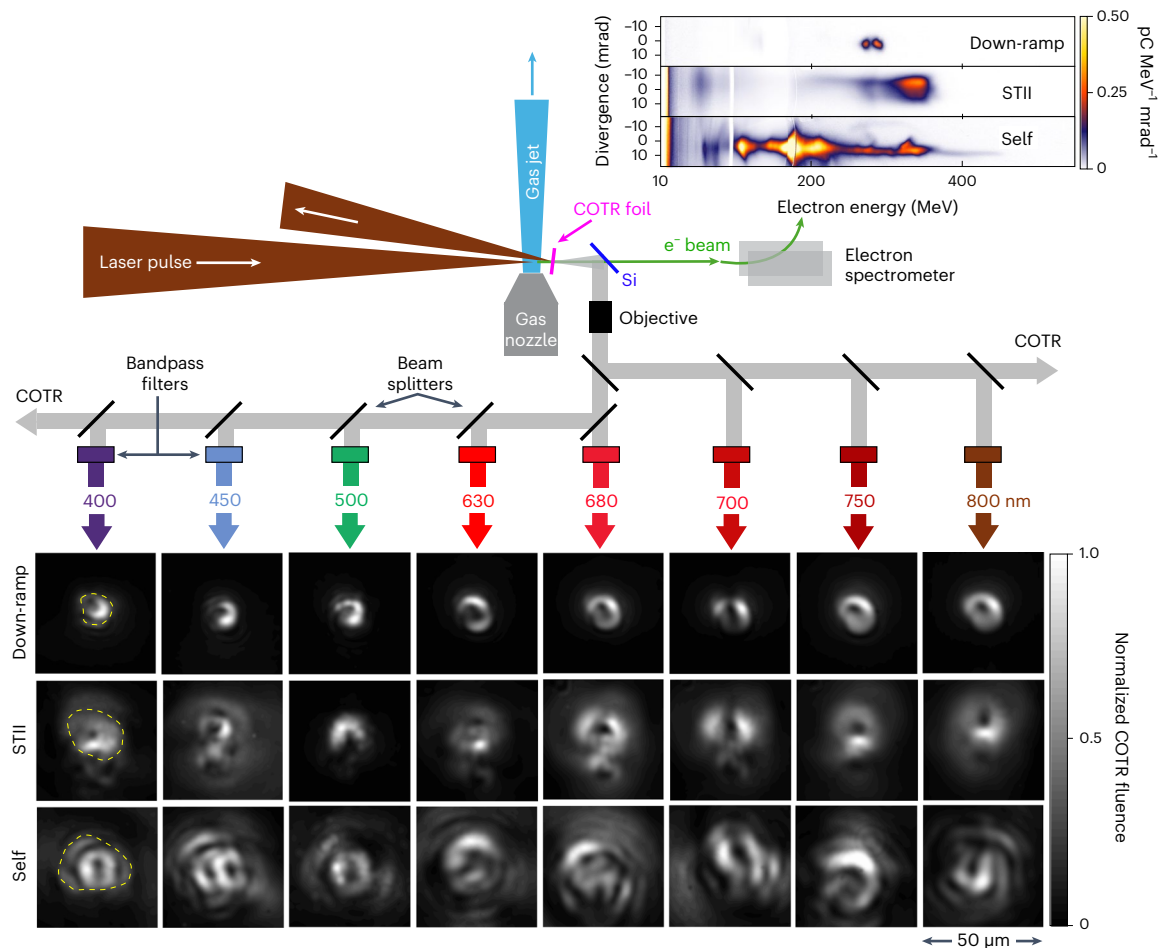


Fig. 1 | Multispectral COTR imaging experiments. Experimental setup (top) and representative electron spectra (top right) for three injection schemes. The total charge above 100 MeV was 54 pC (down-ramp), 280 pC (STII) and 865 pC (self-injection). The laser pulses are focused to a 20 μm spot in the 3-mm-long helium gas jet, creating a plasma of electron density $n_e = 2.3 \times 10^{18} \text{ cm}^{-3}$ (down-ramp and STII) or $4 \times 10^{18} \text{ cm}^{-3}$ (self-injection) and a nonlinear wake. For down-ramp injection, a knife edge perturbed the gas flow, creating a shock near the jet

entrance; for STII, the gas was doped with 1% nitrogen. Representative single-shot, eight-panel COTR image sets (bottom) from three types of injection (rows), bandpass filtered (10 nm FWHM) at eight wavelengths (columns), acquired on the same shots as the electron spectra in the top right. The yellow dashed lines on 400 nm images exemplify the method for defining the area of each pattern. Each image uses a grey scale that optimizes the visibility of its details.

$N_e = N_c + N_i \approx 10^9$ total electrons as determined from the integrated electron spectrometer³⁸ signal, the COTR dominates even for N_c/N_e as small as $\sim 10^{-4}$. By crudely assuming $\|F(\mathbf{k})\| = 1$, we set a lower limit on N_c/N_e . Figure 2b plots the resulting N_c/N_e value as a function of λ , and shows that $N_c/N_e \ll 1$, although it is higher for down-ramp (blue) and self-injection (green) than for STII (orange). The COTR image analyses shown below will reveal that, in fact, $\|F(\mathbf{k})\| \ll 1$, implying that N_c/N_e approaches 1.

Synthetic COTR images. As a prelude to reconstructing unknown e^- bunches, the leftmost column in Fig. 3 shows synthetic (that is, simulated) e^- bunches with four different internal structures: pancake shaped, with Gaussian lengths $\sigma_z = 0.3$ and $0.4 \mu\text{m}$ (rows 1 and 2, respectively); two beamlets (row 3); and a sine wave (row 4). The remaining columns show the forward-calculated images of COTR (Methods) that each bunch generated at each of the four wavelengths. The images strikingly vary in shape, intensity and λ dependence, illustrating the technique's sensitivity to sub-micrometre beam morphology.

The bunch in row 1 is shorter than all the observation wavelengths, and yields an annular COTR profile at all λ values. The central minimum results from the radial polarization of the COTR: opposed polarizations converge and destructively interfere on the z axis in the image plane. The longer similarly shaped bunch (row 2) generates markedly less

COTR near 400 nm, because emissions from its front and rear now destructively interfere. As an analogy, optical thin-film interference patterns display similar sensitivity to film-thickness changes of only $\lambda/4$, which interchange destructive and constructive interferences from the front and back surfaces of the film, respectively. The interference of COTR from the two e^- beamlets in row 3 creates the analogue of a two-slit interference pattern. With 400 nm filtering, oppositely directed electric fields from the two beamlets overlap and interfere destructively on the z axis to produce an axial minimum. As λ lengthens, this feature evolves towards constructive interference and an axial maximum at 800 nm. The sinusoidal electron distribution of the bunch in row 4 yields an axial COTR minimum (maximum) at $\lambda = 400$ (800) nm as in the previous example, but with very different off-axis features. The varied structure of the COTR images suggests the possibility of uniquely reconstructing the internal coherently radiating microstructure of LPA e^- bunches from them. Of course, simple geometric transformations (for example, reflections about the $z = 0$ plane) of asymmetric electron distributions such as rows 3 and 4 in Fig. 3 can yield identical COTR image sets, making them inherently indistinguishable by COTR analysis. Apart from such unavoidable redundancy, the following analysis of experimental COTR images explores to what extent, with what assumptions and with what uncertainty this inverse problem can be uniquely solved.

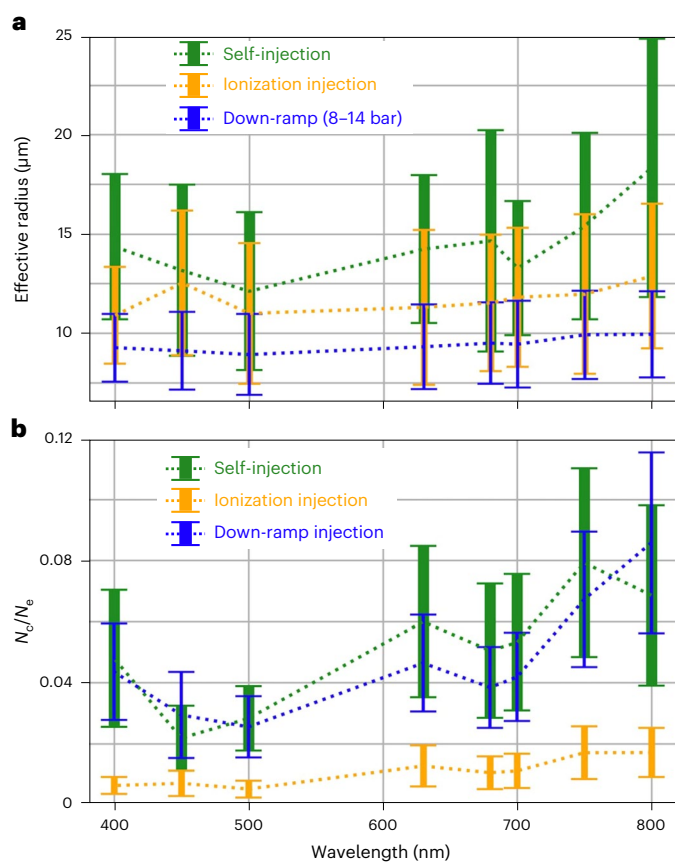


Fig. 2 | Empirical trends in multispectral COTR data. a,b, Effective radii $\sqrt{A/\pi}$ of the COTR radiation patterns (a) and ratio (N_c/N_e) of the number of coherently radiating electrons N_c to the total number of electrons over 100 MeV N_e assuming form factor $F(\mathbf{k}) = 1$ (b), both plotted versus λ for LWFA electrons injected by three methods. Average N_c of 0.3 , 1.8 and 5.4×10^9 for down-ramp, STII and self-injection, respectively. Data are presented as mean values \pm standard deviation over 100 down-ramp, 65 STII and 30 self-injected shots.

Substructure of wakefield-accelerated e^- bunches

Down-ramp-injected e^- bunches. Further progress in unfolding a 3D coherently radiating e^- -bunch substructure requires a quantitative analysis of the multispectral COTR image sets. Down-ramp-injected e^- bunches are a good starting point because of the high coherence, simple structure and near- λ independence of their COTR images, hinting at a simple underlying bunched structure. However, these image sets alone do not adequately constrain the wide range of possible longitudinal structures that iterative reconstructions must sample in search of a 3D solution. Fortunately, the longitudinal profile of a down-ramp-injected bunch can be estimated from its energy spectrum if, as here, three conditions are met: (1) the injected charge is small enough (here 30 pC) to avoid beam loading the wake (which requires ~ 300 pC (ref. 35)), so that the bunch chirps linearly as it accelerates^{39,40}; (2) the down-ramp—typically a shock wave³⁰—is short and steep enough to deterministically initiate both acceleration and chirping at a precise location in the gas jet; (3) the drive pulse focus is well matched to the plasma-guiding forces such that the wake evolves quasi-statically during acceleration. If these conditions are met, then the bunch's energy spectrum maps linearly onto its longitudinal profile (see the 'Calculations and simulations of wakefield acceleration' section). Figure 4a (blue curve) shows a longitudinal profile thus estimated from the spectrum in Fig. 1 (top right). Its ~ 1 μm full-width at half-maximum (FWHM) and 10% energy spread are consistent with our own (Supplementary Fig. 1c) and published⁴¹ simulations of similar accelerators. Moreover, its two internal peaks, which correspond to two peaks in the

energy spectrum, are shorter than the detected wavelengths, leading (Fig. 3, row 1) to coherent, λ -independent annular COTR images. Other methods that do not rely on conditions (1)–(3) above (for example, multi-octave COTR spectroscopy²⁷) can also provide the initial longitudinal profile estimate.

We then used a differential evolution algorithm⁴² to iteratively reconstruct the 3D profiles. Briefly, at each step n , we forward calculated the multispectral COTR image sets from candidate electron distributions $\rho_n(x, y, z)$, evaluated residuals, then updated ρ_n to ρ_{n+1} , guided by a cost function that favoured evolution towards the estimated longitudinal profile, the measured total charge q_m , and an improved match to the COTR data. The algorithm, however, made no assumptions about the symmetry or longitudinal–transverse separability of candidate ρ_n . Each reconstruction j converged towards a solution ρ_j that minimized the cost function. We then repeated the reconstructions dozens of times with randomly initialized $\rho_{n=1}$. We evaluated the set $\{\rho_j\}$ of converged solutions for each shot for determining the quality of fit to the COTR data and for uniqueness, that is, clustering around a single centroid solution, or a small group of geometrically related solutions. Methods provides further details, and Supplementary Fig. 2 shows the reconstruction algorithm.

Figure 4b shows an average of several dozen 3D reconstructions (green) based on the down-ramp-injected COTR data in Fig. 1. The two z -separated sub-bunches described above survive. However, in the course of optimizing its fit to the COTR data, the algorithm modified their transverse and longitudinal shapes from the initial estimate (Fig. 4a, orange curve). This distribution is neither symmetric nor separable. Figure 4c directly compares the measured (top) and reconstructed (bottom) COTR images, illustrating the high level of detail that the reconstruction captured across eight detected λ . Normalized sum-of-squares (NSS) errors (Methods), plotted versus λ in Fig. 4e, were similar in magnitude for individual reconstructions (grey) as for the average of each of the three near-equally populated clusters of reconstructions that principal-component analysis (PCA) identified, as displayed in the inset (Methods provides the details of PCA). The NSS equivalence of individual and cluster-averaged structures is a strong indicator of intracluster uniqueness, since all the cluster members must share common structural attributes to survive averaging. Moreover, the three clusters prove to be close geometric relatives of each other, rather than independent solutions of the inverse problem (Supplementary Fig. 4). Thus, apart from this unavoidable geometric ambiguity, all of the 120 reconstructions clustered tightly around three otherwise equivalent solutions. Supplementary Figs. 3 and 5 show examples of the analyses of other shots.

To re-evaluate $F(\mathbf{k})$ and N_c/N_e , we ran reconstructions that included 90, 80, 70, 60 and 50 per cent of the total charge. NSS increased monotonically as this percentage decreased, as shown by the coloured curves in Fig. 4d. From this, we infer that nearly all of the measured charges contributed to the COTR to some extent, that is, $N_c/N_e \approx 1$, much higher than numbers obtained by assuming $F(\mathbf{k}) = 1$ (Fig. 2b). This, in turn, implies $F(\mathbf{k}) < 1$. Evidently, transverse coherence is small, since the transverse FWHM of the reconstructed down-ramp-injected bunches averages ~ 6 μm , much larger than the detected λ . This re-affirms the importance of analysing both transverse and longitudinal beam structures.

Supplementary Figs. 3–5 illustrate the analyses of other down-ramp-injected bunches, which yield similarly sized envelopes with widely varying internal structures. Envelope reproducibility is a hallmark of down-ramp injection from short, steep, stable shocks. But these results show that the sub-micrometre internal structure of such bunches is correspondingly stochastic. Evidently, fine-grained non-uniformities in nonlinear shocks imprint themselves on injected bunches in ways that are difficult to control, diagnose or simulate. Moreover, off-axis electrons, which down-ramps inject plentifully, evolve unpredictably during acceleration. This internal stochasticity could influence, for example, gain in LWFA-driven FELs.

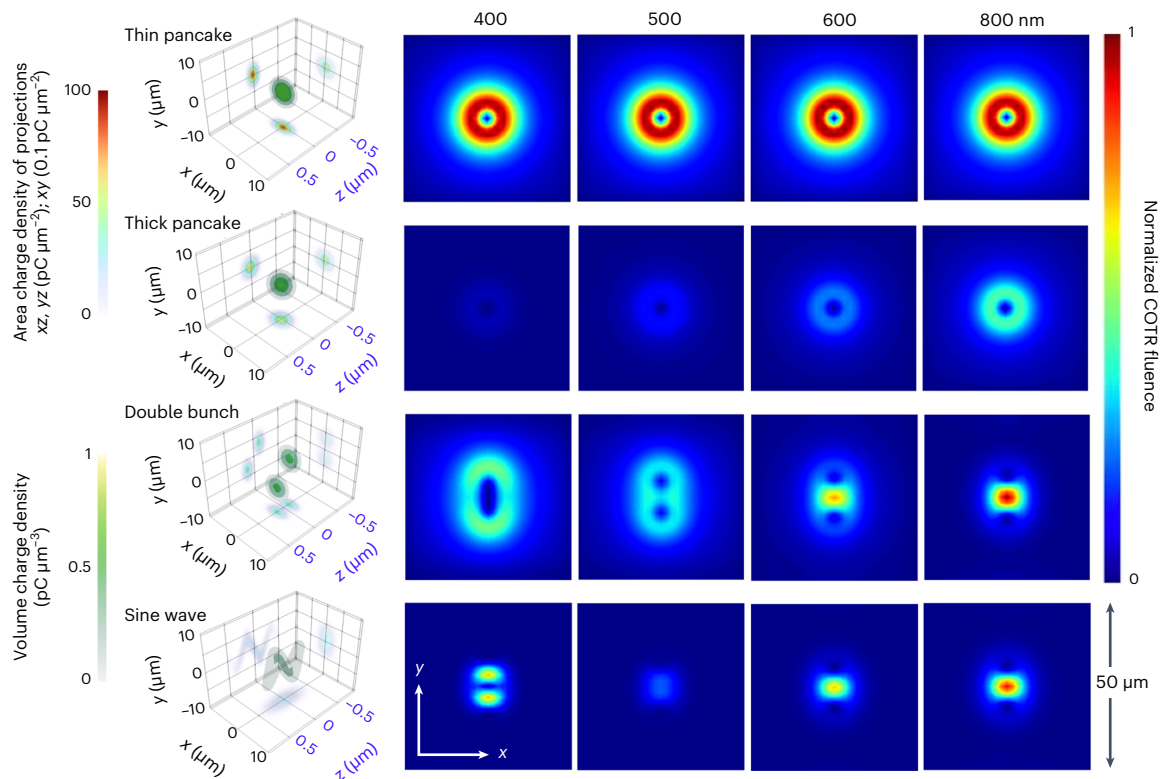


Fig. 3 | Multispectral COTR images of synthetic electron distributions.

The left column shows the four electron distributions: width $\sigma_{\perp} = 3 \mu\text{m}$, length $\sigma_z = 300 \text{ nm}$ (row 1) and 400 nm (row 2) FWHM longitudinally Gaussian distributions; two 300 nm FWHM longitudinal Gaussians separated by 400 nm longitudinally and $6 \mu\text{m}$ transversely (row 3); a distribution with sinusoidal displacements along y and 800 nm period along z (row 4). Two-dimensional

projections of each distribution are shown on the x - y , x - z and y - z planes. The total charge ratio is 1.0:1.0:1.0:1.1; electron energy, 285 MeV throughout. The remaining rows show the COTR images generated from the corresponding electron distribution at $400, 500, 600$ and 800 nm with a collection angle of 0.28 mrad . The colour scale is consistent in each column, and normalized to the corresponding image in the first row.

Ionization-injected e^- bunches. STII occurs over an extended region of the gas jet, and generates e^- bunches of sufficient charge ($\sim 300 \text{ pC}$) to load the wake³⁵. Thus, the method used above to estimate the longitudinal profile of down-ramp-injected e^- bunches is not valid. Fortunately, other information is available. A recent multioctave COTR spectrometry study of STII bunches from DRACO yielded a bunch length of $\sigma_z \approx 3 \mu\text{m}$ (ref. 27). The PIC-simulated STII e^- bunch (Fig. 5a), generated under our experimental conditions and propagated to the location of the COTR foil, corroborates this measurement. Moreover, it reveals two concentrations of optical-wavelength density displacements that are candidate COTR sources: (1) oscillations with $\lambda \approx 800 \text{ nm}$ in the leading portion of the beam, from interaction of the e^- bunch with the trailing edge of drive laser pulse; (2) shorter-wavelength oscillations in the trailing portion, from late off-axis ionization injection from N^{7+} (ref. 43). Figure 5b shows the extracted parameterized model of the first source that matches its simulated amplitude, location and frequency spectrum, but enables fine variations to fit the data (Methods). Figure 5c shows a representation of the second source that emerged from such a fit. Figure 5d shows them combined with the featureless intervening section, which generated a negligible COTR.

Figure 5e (top row) shows the COTR data from a typical STII beam. The second row shows the calculated COTR from the source in Fig. 5b. When iteratively optimized, it reproduced the salient features of longer- λ images, but failed to model those at shorter λ . Conversely, the source in Fig. 5c reproduced the main features of short- λ images, but contributed negligibly at $\lambda > 630 \text{ nm}$. COTR from the optimized combined sources (Fig. 5e, last row), however, reproduced the main observed features at all λ . NSS was no larger than that for the

down-ramp-injected bunches, and with repeated reconstructions of one shot or different shots, the oscillatory leading and trailing features consistently reappeared. Only the detailed internal structure of the latter varied widely among the reconstructions of even one shot and thus could not be uniquely determined.

Self-injected e^- bunches. Attempted reconstructions of self-injected e^- bunches led to widely divergent solutions with high NSS. Supplementary Figs. 7 and 8 show a representative example.

Discussion

The solution to the inverse problem of reconstructing $\rho(x, y, z)$ from COTR intensity measurements is inherently uncertain because the optical-phase information that links ρ to the COTR field $E_{\perp}(x, y)$ is lost. Nevertheless, similarly handicapped image reconstruction methods are used with high confidence in, for example, medicine and structural engineering. Success relies on effectively compensating for the missing phase information, screening trial solutions for fit quality and uniqueness, and repeatability of the iterative procedure. Here independent measurements of electron charge and energy, knowledge that bunch charge is negative, and reliable guidance on longitudinal bunched structure from electron spectra or PIC simulations provided compensating information that constrained possible solutions, and is readily available to most LWFA researchers. A high-quality fit of a calculated COTR image set to a measured one is a necessary, but not sufficient, condition for trusting a trial solution. Here we quantified the fit quality by calculating the NSS values based on a pixel-to-pixel comparison. Visually satisfying fits such as those shown in Figs. 4 and 5 for the down-ramp-injected and STII bunches, respectively, generated

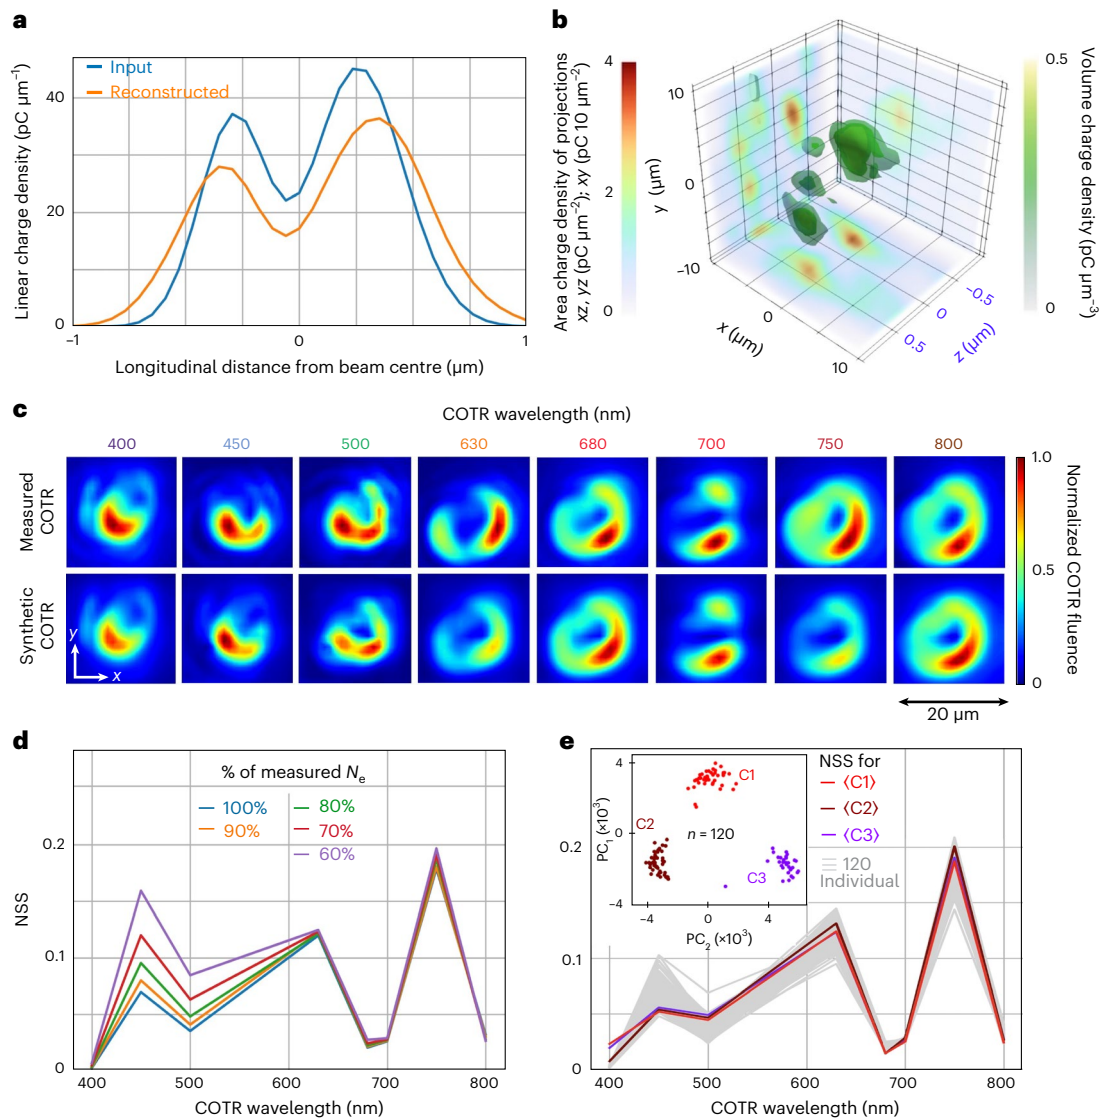


Fig. 4 | 3D reconstruction of down-ramp-injected e^- bunch. **a**, Input longitudinal profile of the e^- bunch derived from its electron spectrum (blue), and its average profile reconstructed from the COTR data (orange). **b**, Average of 47 reconstructions from one cluster of solutions for the 3D electron density of the e^- bunch, with average two-dimensional projections plotted on the bounding planes. The beam is much wider than it is long, so the z axis is scaled by a factor of 10. **c**, The top row shows the measured COTR at eight wavelengths. The bottom row shows the forward-calculated COTR from the cluster-averaged structure in **b**. The colour scale is different at each λ , but the same for the measured and

reconstructed COTR at each λ . **d**, Average NSS at each λ for 120 reconstructions that included 100% (blue), 90% (orange), 80% (green), 70% (red) and 60% (purple) of the measured total charge. **e**, The grey area shows the superposed NSS values for 120 individual reconstructions, using 100% of the measured charge. The inset shows a PCA plot, revealing three colour-coded clusters of reconstructions (axes: the first two principal components of the distribution of reconstructions). The coloured plots in the main panel, NSS for structures averaged over each correspondingly colour-coded cluster. Reconstruction in **b** belongs to the red cluster.

NSS values comparable with those achieved with synthetic data such as those shown in Fig. 3.

A trustworthy $\rho(x, y, z)$ must also be unique. Here, to quantify uniqueness, we repeatedly ran the reconstruction algorithm on each dataset with randomly varied initial parameters to evaluate solution clustering⁴⁴. For successful reconstructions, solutions grouped into two or three geometrically related clusters, in each of which the cluster-averaged structure yielded NSS as small as—or nearly as small as—that of individual reconstructions, demonstrating uniqueness. The difficulty of reconstructing self-injected e^- bunches reflected lack of guidance on their longitudinal structure, their wide energy spread and the complexity of the COTR patterns.

Several extensions of multispectral COTR imaging can probably increase its diagnostic capability in the future. First, there is enough COTR energy to distribute to almost three times as many filtered

cameras as demonstrated here, enabling larger spectral range and density and correspondingly greater resolution. Extensions to shorter λ will be important for diagnosing nanobunching in LWFA-driven XFELs^{45,46} and nanoprebunching schemes¹⁹. Second, by conveying the e^- bunch after the COTR imaging split-off mirror (Fig. 1, top) to a multioctave spectrometer, one can combine the latter's wide spectral range with the fine lateral resolution of multispectral imaging. Third, e^- -beam perturbations by laser-block and COTR foils, though carefully minimized here, can be eliminated by using apertured^{47,48} foils. For applications remote from the LWFA, a laser block is not needed.

In conclusion, we have demonstrated multispectral COTR imaging as a versatile, low-tech, high-resolution diagnostic of the internal coherently radiating structure of LWFA e^- bunches that are quasi-monoenergetic and for which the longitudinal charge profile can be reliably estimated. Using this tool, we have identified optical

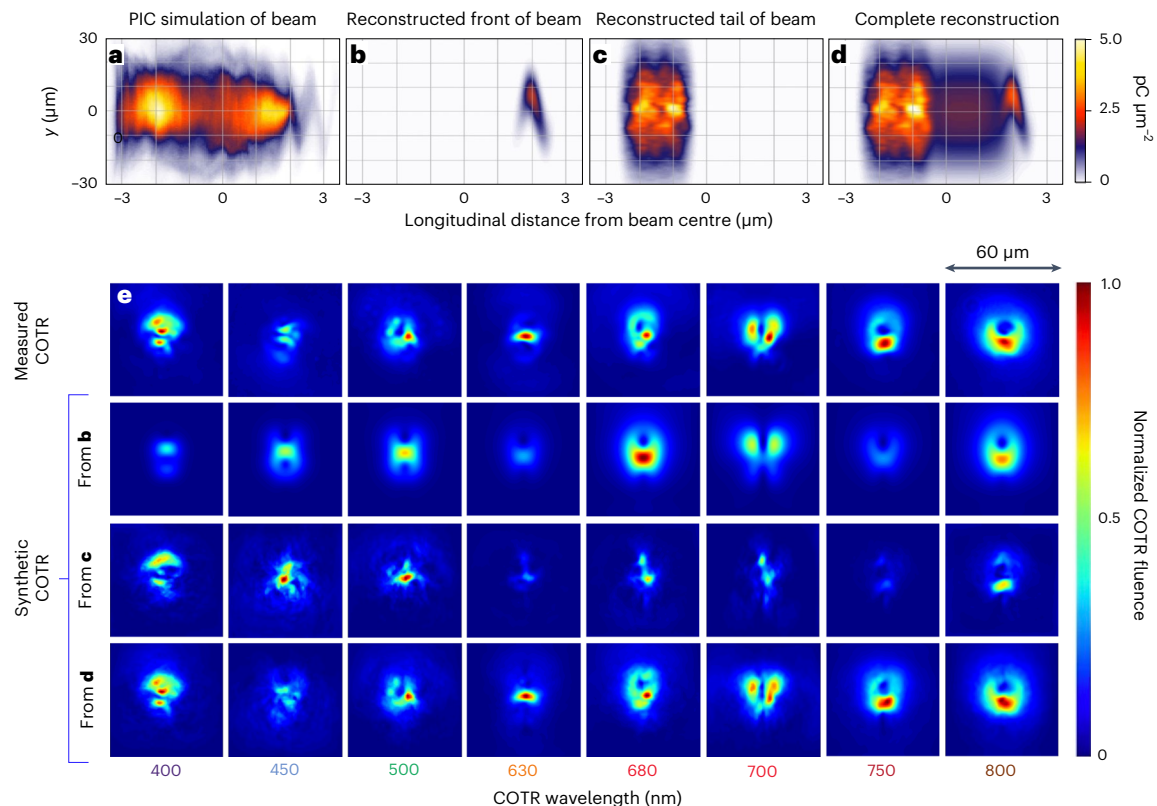


Fig. 5 | 3D representation of ionization-injected e^- bunch. **a**, PIC simulation of density distribution $\rho(x, y, z)$ of the STII e^- bunch after exiting the LWFA, projected onto the y - z plane. **b, c**, Parameterized representation of laser-induced modulation at the leading edge (**b**) and ionization-induced higher-frequency modulation at the tail of **a** (**c**). **d**, Combination of the data in **b** and **c** with the

unmodulated middle section of the data in **a**, **e**. The first row shows the measured COTR images from a single shot at the eight corresponding central wavelengths. The bottom three rows show the best-fit COTR images generated by the optimized distributions in **b–d**. The colour scale in each column is normalized to the brightest image in that column.

and structural fingerprints of down-ramp- and ionization-injected e^- beams, and reconstructed the internal fine structure with unprecedented detail. We anticipate that this approach will play a key role in diagnosing micro- and nanobunching in future tabletop LWFA-driven X-ray sources, as well as in optimizing their performance through machine learning strategies.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41566-024-01475-2>.

References

- Pellegrini, C., Marinelli, A. & Reiche, S. The physics of X-ray free-electron lasers. *Rev. Mod. Phys.* **88**, 015006 (2016).
- Bergman, U., Yachandra, V. K. & Yano, J. *X-Ray Free Electron Lasers: Applications in Materials, Chemistry and Biology* Vol. 18 (Royal Society of Chemistry, 2017).
- Bonifacio, R., Pellegrini, C. & Narducci, L. M. Collective instabilities and high-gain regime free electron laser. *AIP Conf. Proc.* **236**, 236–259 (1984).
- Gover, A. et al. Superradiant and stimulated-superradiant emission of bunched electron beams. *Rev. Mod. Phys.* **91**, 35003 (2019).
- Frank, I. M. & Ginzburg, V. L. Radiation of a uniform moving electron due to its transition from one medium into another. *J. Phys. (USSR)* **9**, 353–362 (1945).
- Schroeder, C. B., Esarey, E., van Tilborg, J. & Leemans, W. P. Theory of coherent transition radiation generated at a plasma-vacuum interface. *Phys. Rev. E* **69**, 016501 (2004).
- Rosenzweig, J., Travish, G. & Tremaine, A. Coherent transition radiation diagnosis of electron beam microbunching. *Nucl. Instrum. Methods Phys. Res., Sect. A* **365**, 255–259 (1995).
- Tremaine, A. et al. Observation of self-amplified spontaneous-emission-induced electron-beam microbunching using coherent transition radiation. *Phys. Rev. Lett.* **81**, 5816–5819 (1998).
- Lumpkin, A. H. et al. Evidence for microbunching ‘sidebands’ in a saturated free-electron laser using coherent optical transition radiation. *Phys. Rev. Lett.* **88**, 234801 (2002).
- Labat, M. et al. Seeded free-electron laser driven by a compact laser plasma accelerator. *Nat. Phys.* **17**, 150–156 (2022).
- Wang, W. et al. Free-electron lasing at 27 nanometres based on a laser wakefield accelerator. *Nature* **595**, 516–520 (2021).
- Pompili, R. et al. Free-electron lasing with compact beam-driven plasma wakefield accelerator. *Nature* **605**, 659–662 (2022).
- Galletti, M. et al. Stable operation of a free-electron laser driven by a plasma accelerator. *Phys. Rev. Lett.* **129**, 234801 (2022).
- Nakajima, K. Towards a table-top free-electron laser. *Nat. Phys.* **4**, 92–93 (2008).
- Steiniger, K. et al. Building on optical free-electron lasers in the traveling-wave Thomson-scattering geometry. *Front. Phys.* **6**, 00155 (2019).
- White, G. R. & Raubenheimer, T. O. Transverse jitter tolerance issues for beam-driven plasma accelerators. In *Proc. 10th International Particle Accelerator Conference (IPAC'19)* 3774–3777 (JACoW Publishing, 2019).

17. Xu, X. et al. Nanoscale electron bunching in laser-triggered ionization injection in plasma accelerators. *Phys. Rev. Lett.* **117**, 034801 (2016).
18. Lumpkin, A. H. et al. Coherent optical signatures of electron microbunching in laser-driven plasma accelerators. *Phys. Rev. Lett.* **125**, 014801 (2020).
19. Xu, X. et al. Generation of ultrahigh-brightness pre-bunched beams from a plasma cathode for X-ray free-electron lasers. *Nat. Commun.* **13**, 3364 (2022).
20. Grüner, F. et al. Design considerations for table-top, laser-based VUV and X-ray free electron lasers. *Appl. Phys. B* **86**, 431–435 (2007).
21. Steiniger, K. et al. Optical free-electron lasers with traveling-wave Thomson-scattering. *J. Phys. B: At. Mol. Opt. Phys.* **47**, 234011 (2014).
22. Emma, C. et al. Terawatt attosecond X-ray source driven by a plasma accelerator. *APL Photonics* **6**, 076107 (2021).
23. Downer, M. C., Zgadaj, R., Debus, A., Schramm, U. & Kaluza, M. C. Diagnostics for plasma-based electron accelerators. *Rev. Mod. Phys.* **90**, 035002 (2018).
24. Glinec, Y., Faure, J., Norlin, A., Pukhov, A. & Malka, V. Observation of fine structures in laser-driven electron beams using coherent transition radiation. *Phys. Rev. Lett.* **98**, 98–101 (2007).
25. Bajlekov, S. I. et al. Longitudinal electron bunch profile reconstruction by performing phase retrieval on coherent transition radiation spectra. *Phys. Rev. ST Accel. Beams* **16**, 040701 (2013).
26. Heigoldt, M. et al. Temporal evolution of longitudinal bunch profile in a laser wakefield accelerator. *Phys. Rev. ST Accel. Beams* **18**, 121302 (2015).
27. Zarini, O. et al. Multioctave high-dynamic range optical spectrometer for single-pulse, longitudinal characterization of ultrashort electron bunches. *Phys. Rev. Accel. Beams* **25**, 012801 (2022).
28. Bakkali Taheri, F. et al. Electron bunch profile reconstruction based on phase-constrained iterative algorithm. *Phys. Rev. Accel. Beams* **19**, 032801 (2016).
29. Liu, Y. et al. Experimental observation of femtosecond electron beam microbunching by inverse free-electron-laser acceleration. *Phys. Rev. Lett.* **80**, 4418–4421 (1998).
30. Schmid, K. et al. Density-transition-based electron injector for laser driven wakefield accelerators. *Phys. Rev. ST Accel. Beams* **13**, 091301 (2010).
31. Mirzaie, M. et al. Demonstration of self-truncated ionization injection for GeV electron beams. *Sci. Rep.* **5**, 14659 (2015).
32. Scheinker, A., Cropp, F., Paiguaga, S. & Filippetto, D. An adaptive approach to machine learning for compact particle accelerators. *Sci. Rep.* **11**, 19187 (2021).
33. Rodenburg, J. M. et al. Hard-X-ray lensless imaging of extended objects. *Phys. Rev. Lett.* **98**, 034801 (2007).
34. Marchesini, S. et al. Coherent X-ray diffractive imaging: applications and limitations. *Opt. Express* **11**, 2344–2353 (2003).
35. Couperus, J. P. et al. Demonstration of a beam loaded nanocoulomb-class laser wakefield accelerator. *Nat. Commun.* **8**, 487 (2017).
36. Raj, G. et al. Probing ultrafast magnetic-field generation by current filamentation instability in femtosecond relativistic laser-matter interactions. *Phys. Rev. Research* **2**, 023123 (2020).
37. Hannasch, A. et al. Nonlinear inverse Compton scattering from a laser wakefield accelerator and plasma mirror. Preprint at <https://arxiv.org/abs/2107.00139> (2021).
38. Kurz, T. et al. Calibration and cross-laboratory implementation of scintillating screens for electron bunch charge determination. *Rev. Sci. Instrum.* **89**, 093303 (2018).
39. Kostyukov, I., Pukov, A. & Kiselev, S. Phenomenological theory of laser-plasma interaction in ‘bubble’ regime. *Phys. Plasmas* **11**, 5256–5264 (2004).
40. Lin, C. et al. Long-range persistence of femtosecond modulations on laser-plasma-accelerated electron beams. *Phys. Rev. Lett.* **108**, 094801 (2012).
41. Ke, L. T. et al. Near-GeV electron beams at a few per-mille level from a laser wakefield accelerator via density-tailored plasma. *Phys. Rev. Lett.* **126**, 214801 (2021).
42. Storn, R. & Price, K. Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *J. Global Optim.* **11**, 341–359 (1997).
43. Köhler, A. et al. Restoring betatron phase coherence in a beam-loaded laser-wakefield accelerator. *Phys. Rev. Accel. Beams* **24**, 091302 (2021).
44. Everitt, B. S., Landau, S., Leese, M. & Stahl, D. *Cluster Analysis* 5th edn (Wiley, 2011).
45. Gazazian, E., Ispirian, K., Ispirian, R. & Ivanian, M. Measurement of very short time 10^{-19} to 10^{-17} s structures with the help of X-ray transition radiation. *Nucl. Instrum. Methods in Phys. Res., Sect. B* **173**, 160–169 (2001).
46. Lumpkin, A., Fawley, W. & Rule, D. A concept for Z-dependent microbunching measurements with coherent X-ray transition radiation in a SASE FEL. In *Proc. 26th International Free Electron Laser Conference and 11th FEL Users Workshop* 515–518 (2004).
47. Karlovets, D. & Potylitsyn, A. On the theory of diffraction radiation. *J. Exp. Theor. Phys.* **107**, 755–768 (2008).
48. Lumpkin, A. H., Berg, W. J., Dooling, J., Sun, Y., Wootton, K. P., Rule, D. W., Murokh, A. & Musumeci, P. Proposed research with microbunched beams at LEA. In *10th Int. Beam Instrum. Conf. (IBIC2021)* 244–248 (JaCoW Publishing, 2021).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024

Methods

Laser wakefield electron acceleration

DRACO⁴⁹ at HZDR drove plasma wakes with laser pulses of 30 fs duration, 800 nm centre wavelength and -2 J energy. An $f/20$ off-axis parabolic mirror with 2 m focal length focused these pulses to a spot size of 20 μm (FWHM) onto the entrance plane of a 3-mm-long Mach-10 helium (He) gas jet. The leading edge of the laser pulse fully ionized the He, after which its intense peak drove an LWFA in the He plasma. For down-ramp and ionization injection experiments, the plasma in the jet's central density plateau (length $L \approx 1.5$ mm) had an average electron density in the range of $2.0 \times 10^{18} < n_e < 3.3 \times 10^{18} \text{ cm}^{-3}$. For self-injection experiments, it had an average density of $n_e \approx 4 \times 10^{18} \text{ cm}^{-3}$. For down-ramp injection, a knife edge was inserted into the gas jet, creating a thin dense shock near the beginning of the plateau, pinpointed by a transverse wide-bandwidth shadowgraphy probe with a centre wavelength of 800 nm (ref. 50) (Supplementary Fig. 1b). Laser focus and shock positions were scanned to optimize the LWFA stability and performance. For STII, we removed the shock and doped the gas with 1% nitrogen. We then scanned the pulse shape with a programmable acousto-optic dispersive filter (Fastlite Dazzler) and adjusted the pulse energy to optimize the LWFA performance. For self-injection, we used a pure He jet with no shock and shifted the laser focus several millimetres before the jet. This mismatched the laser divergence and the plasma wake's focusing force, inducing radial wake oscillations that triggered the injection of plasma electrons into the wake.

A magnetic electron spectrometer with its entrance plane at $z = 30$ cm downstream of the gas-jet exit determined the electron energy distribution for each shot⁴⁹. A Konica Minolta OG 400 scintillating screen recorded the dispersed electron beam (Fig. 1 (top right) shows examples). We converted the luminescence intensity into charge per unit energy per pixel using methods described elsewhere³⁸. The absolute charge calibration uncertainty was ~20%. The r.m.s. shot-to-shot charge fluctuations were only a few per cent. Uncertainties in the >200 MeV electron energy measurement primarily originated from pointing and divergence fluctuations of LWFA electrons entering the spectrometer and amounted to ~2% for electrons in the range of 200–350 MeV.

Calculations and simulations of wakefield acceleration

In the 'Substructure of wakefield-accelerated e^- bunches' section, we estimated the longitudinal profile $N_e(\xi)$ and length $\Delta\xi$ of down-ramp-injected e^- bunches from their measured energy distribution $\partial N_e/\partial U_e$ and energy spread ΔU_e , respectively. Here $\xi = z - v_0 t$ is the longitudinal coordinate in the frame of a bubble moving with velocity v_0 along the lab coordinate z . This estimate assumed a quasi-static bubble and an accelerating field $E_z = (en_e/2\epsilon_0)\xi$ that is linear in ξ but exceeds the field $(en_e/3\epsilon_0)\xi$ of a uniformly charged sphere because of the concentration of electrons at its rear³⁹. The PIC simulations of down-ramp-injected LWFA for our conditions, such as the example shown in Supplementary Fig. 1c, validate these assumptions for our conditions. Linear E_z is expected when accelerated charge, which averaged -50 pC for our down-ramp-injected LWFA, is well below the beam-loading limit of -300 pC (ref. 35). The estimate also assumed that injection is localized in space and time. Transverse shadowgraphs⁵⁰ such as the example shown in Supplementary Fig. 1b validated this assumption. Under these conditions, an electron injected at position ξ_i within the bubble and accelerated to ξ_f , positions that are at potentials $\Phi_x = (en_e/4\epsilon_0)\xi_x^2$ (where $x = i, f$) with respect to the bubble centre, gains energy $U_e = e(\Phi_i - \Phi_f)$ in the bubble frame. Applying the chain rule $dN_e/d\xi = (\partial N_e/\partial U_e)(\partial U_e/\partial \xi)$ and transforming to the lab frame, we find that $\partial N_e/\partial U_e$ maps onto the bunch's longitudinal profile dN_e/dz via the linear chirp $\partial U_e/\partial z$, that is,

$$\frac{dN_e}{dz} = \frac{n_e e^2}{2\epsilon_0} \frac{\partial N_e}{\partial U_e} L_{\text{acc}}, \quad (1)$$

where L_{acc} is the acceleration length in the lab frame. This blue curve in Fig. 4a plots this function using $L_{\text{acc}} \approx L$, $n_e = 2.3 \times 10^{18} \text{ cm}^{-3}$ and the measured $\partial N_e/\partial U_e$ (Fig. 1, top right). For the down-ramp-injected shot analysed in the main text, the chirp was $\partial U_e/\partial \xi \approx 3 \text{ MeV } \mu\text{m}^{-1}$ and bunch length $\Delta\xi = (2\epsilon_0/e^2 n_e)(\Delta U_e/L) \approx 850 \text{ nm}$ FWHM for -10% energy spread at 280 MeV, which is consistent with published⁴¹ and our own data (Supplementary Fig. 1c shows the PIC simulations for this injection method). Since our spectrometer has a resolution of $\Delta U_e^{(\text{res})} \approx 10 \text{ MeV}$, this approach determines bunch length $\Delta\xi$ (bunch duration $\Delta\xi/c$) with 300 nm (1 fs) resolution. The r.m.s. fluctuation of ~20% in $n_e L_{\text{acc}}$ was the principal uncertainty in the estimated dN_e/dz and $\Delta\xi$. In iterative reconstructions of the 3D profile $N_e(x, y, \xi)$ of down-ramp-injected bunches, the estimated longitudinal profile served as a parameterized initial guess (rather than a fixed function) that was allowed to vary over a range consistent with this experimental uncertainty, yielding, for example, the orange curve in Fig. 4a.

The 3D PIC simulations of STII were performed using the code PIConGPU⁵¹ to evaluate the spatial structure of the bunch and the longitudinal phase space. For the STII simulation (Fig. 5a), we used a simulation box of $768 \times 4,608 \times 768$ cells with a transverse resolution of 4.5 and a longitudinal resolution of 36 sampling points per laser wavelength. The laser is modelled using a Gauss-Laguerre reconstructed laser profile measured during the experiment. To avoid numerical Cherenkov radiation, the Lehe field solver⁵² was used together with the Boris pusher⁵³ and the Esirkepov current deposition scheme⁵⁴. The exact code version used and all the setup files are available elsewhere⁵⁵.

The COTR images directly calculated from the PIC-simulated density distribution (Fig. 5a) are shown in Supplementary Fig. 6. By extracting the oscillatory leading and trailing edges and smooth middle section of this distribution, we verified that nearly all the calculated COTR data originated from the leading and trailing edges. To optimize the match of the calculated and measured COTR images, we first constructed the following parameterized representation of the leading portion of the Fig. 5a distribution:

$$\rho_1(x, y, z) = A e^{-[(x+\delta x)/2\sigma_x]^2} \times e^{-(z/2\sigma_z)^2} \times e^{\left\{[-y+\delta y - B \cos(\frac{2\pi}{\lambda} z + \phi_z)]/2\sigma_y\right\}^2}, \quad (2)$$

where A is the overall amplitude in units of $\text{C } \mu\text{m}^{-3}$; and x, y and z and all the remaining parameters were normalized to 1 μm : $\sigma_{x,y,z}$ are the Gaussian widths; δx and δy are x and y offsets; and λ and B are the oscillation wavelength and amplitude, respectively. The first two Gaussians in equation (2) constrain the overall distribution in the x and z directions, and the third represents the leading oscillation. The simulated trailing oscillations in Fig. 5a had less well-defined wavelength and amplitude; therefore, we specified only the spatial boundaries of this region and left the reconstruction algorithm to find the density distribution on its own. A smooth super-Gaussian function that generated negligible COTR represented the charge density in the middle section. We then ran dozens of iterative reconstructions with varied starting parameters until a dominant cluster of solutions representing the best fit emerged. Figure 5d represents the best-fit composite bunch distribution, and the bottom row in Fig. 5d shows the corresponding best-fit COTR data. The best-fit equation (2) parameters for the data in Fig. 5 were as follows: $A = 0.323 \pm 0.034 \text{ pC } \mu\text{m}^{-3}$, $B = 7.70 \pm 1.00 \text{ } \mu\text{m}$, $\sigma_x = 3.50 \pm 0.40 \text{ } \mu\text{m}$, $\sigma_y = 5.40 \pm 0.40 \text{ } \mu\text{m}$, $\sigma_z = 0.16 \pm 0.10 \text{ } \mu\text{m}$, $\delta x = -0.65 \pm 0.20 \text{ } \mu\text{m}$, $\delta y = -0.90 \pm 0.06 \text{ } \mu\text{m}$, $\lambda = 0.89 \pm 0.07 \text{ } \mu\text{m}$.

COTR generation, imaging and calibration

A schematic of the 65 μm laser-block foil and 25 μm Kapton COTR foil with a 300-nm-thick aluminium-coated back surface in relation to the gas jet and LWFA drive pulse is shown in another work¹⁸. For the results presented here, the Kapton foil also included a 65 μm low- Z adhesive layer. Here 67 such foil pairs were mounted at the perimeter of a remotely controlled wheel that maintained a fixed

1 mm distance between the foils and rotated a fresh pair into place after each shot. Before the next shot, a lamp temporarily illuminated the new foil's COTR-emitting back surface, enabling it to be finely adjusted into focus. To image the COTR from the foil to detectors, we employed a $\times 10$ infinity-corrected long-working-distance microscope objective (Mitutoyo Model 46-144; numerical aperture, 0.28; focal length, 20 mm) for all the measurements reported here. By inserting a resolution test target in place of the COTR foil, we measured almost $\times 40$ magnification and better than $2 \mu\text{m}$ (σ) resolution at the foil surface. To extend the spectral range to $\lambda = 193 \text{ nm}$, we temporarily used a reflective microscope objective (TECHSPEC 89-723), but detected no signals at this wavelength. All the data reported here were taken with the Mitutoyo objective.

The integrated COTR energy captured at each detector was calibrated by measuring the power of linearly polarized test lasers at 405, 532, 633 and 805 nm at the position of the COTR foil with a power meter and then exposing the cameras for both *s* and *p* polarizations for a known time, yielding a counts per nanojoule calibration. These measurements were checked against the manufacturer's published transmission/reflection curves for beamsplitters and transmission curves for bandpass filters. The COTR is radially polarized, but encountered beamsplitters with polarization- and λ -dependent transmission/reflection coefficients en route to each detector. We strategically designed each optical path to deliver an *s/p* intensity ratio as close to unity as possible. Actual *s/p* ratios ranged from 0.7 to 2.0, except for $\lambda = 700 \text{ nm}$, for which the *s/p* value was ~ 7.0 . Supplementary Table 1 provides a full list. For the data analysis, each image's actual *s/p* ratio was taken into account.

e^- bunch reconstruction procedure

Differential evolution. For reconstructing the normalized 3D charge density $\rho(x, y, z)$ of an e^- bunch from multispectral COTR data, we used a differential evolution⁵⁶ algorithm, a highly parallelizable global-maximum-search algorithm that does not require computationally expensive gradient calculations. At each step '*n*' of an iterative reconstruction, one calculates the COTR field $\mathbf{E}_\perp^{(n)}$ that an e^- bunch of candidate distribution $\rho_n(x, y, z)$ propagating along *z* generates on emerging from a foil in the *x*-*y* plane. This starts with calculating the field produced at transverse displacement \mathbf{r} from the point at which a single electron emerges. After a lens with numerical aperture θ_{max} images it to a detector, this field is⁵⁷

$$\mathbf{E}_\perp^{(\text{PSF})}(\mathbf{r}) = \frac{2ek}{c} \int_0^{\theta_{\text{max}}} \frac{\theta^2 d\theta}{\theta^2 + \gamma^{-2}} J_1(k\theta|\mathbf{r}|) \hat{\mathbf{r}}, \quad (3)$$

where *e* and γ are the electron charge and Lorentz factor, respectively; *k* and *c* are the wave number and speed of emitted transition radiation, respectively; and J_1 is a Bessel function of the first kind. Equation (3) is our point-spread function (PSF). The total transverse COTR field from an e^- bunch is then the convolution of $\rho_n(x, y, z)$, multiplied by the phase delay e^{-ikz} of radiation from each longitudinal slice of ρ_n , with the PSF as follows:

$$\mathbf{E}_\perp^{(n)}(k, x, y) = Q \int dx' \int dy' \mathbf{E}_\perp^{(\text{PSF})}(\mathbf{r} - \mathbf{r}') \int e^{-ikz} \rho_n(x', y', z) dz, \quad (4)$$

where *Q* is the total charge and $|\mathbf{r}'|^2 = x'^2 + y'^2$. The total COTR intensity, which is proportional to $|\mathbf{E}_\perp^{(n)}|^2$, is then compared with the data. At the γ and θ values involved in this work, effects from the finite size of the COTR foil or milliradian-scale divergences⁵⁷ are negligible.

Digital reconstruction. For automated digital reconstruction, we segmented $\rho(x, y, z)$ into N_v voxels with Gaussian longitudinal profiles ($\Delta z_{\text{r.m.s.}} = 80 \text{ nm}$), effectively suppressing radiation at $\lambda < 250 \text{ nm}$ that we never observed²⁷, and transverse widths $\Delta x = \Delta y = 1 \mu\text{m}$, just below

our resolution threshold in the foil plane. The voxel spacing was $\sim 60 \text{ nm}$ longitudinally and $1 \mu\text{m}$ transversely. Thus, $N_v \approx 15,000 = 21 \times 21 \times 35$. The voxel charge was then substituted for a single-electron charge *e* in evaluating the PSF in equation (3), and the convolution integrals in equation (4) became sums over voxels. We similarly segmented the calculated $\mathbf{E}_\perp^{(n)}(k, x, y)$ into pixels for comparison with charge-coupled-device images. For each iteration, a dimensionless NSS error $N_\lambda^{-1} \sum_p (\Delta N_p)^2$ is calculated from differences $\Delta N_p = N_p^{(c)} - N_p^{(m)}$ between the calculated (*c*) and measured (*m*) counts in pixel '*p*' normalized to the summed squared count $N_\lambda = \sum_p (N_p^{(m)})^2$ at each λ . Procedures for aligning the calculated and measured intensity distributions and for evaluating the alignment uncertainties are discussed in ref. 58 and in Supplementary Section 2. We generated a composite NSS for each multi- λ image set.

The overall cost function that reconstructions minimized was a weighted sum of this composite NSS and four additional terms. Three of the latter ensured that the reconstructed $\{\rho_{ij}\}$ (where *i* and *j* are voxel and reconstruction labels, respectively, and brackets denote the set of values for N_v voxels) adhered within designated uncertainties to (1) a prescribed longitudinal profile from an independent measurement or simulation, (2) the total measured charge and (3) a positive-definite electron density ρ . The fourth ensured that it generated (4) the measured overall COTR intensity at each λ within the detector calibration uncertainty. The longitudinal profile term had the form $\sum_\zeta (\Delta_\zeta/q_\zeta)^2 \Theta(\Delta_\zeta)$, where q_ζ is the prescribed charge at longitudinal position ζ , Δ_ζ is the difference $Q_\zeta - q_\zeta$ between the candidate profile's charge Q_ζ at ζ and q_ζ and Θ is the Heaviside step function. Thus, if $Q_\zeta > q_\zeta$, $\Theta(\Delta_\zeta) = 1$ and Q_ζ increased the cost, whereas if $Q_\zeta < q_\zeta$, $\Theta(\Delta_\zeta) = 0$ and Q_ζ incurred no cost. Similarly, the total charge term had the form $\Delta Q^2 \Theta(\Delta Q)$, where ΔQ is the difference $Q - q_m$ between the candidate's *Q* and the measured q_m total charges. Although this term encouraged *Q* to converge towards q_m , we typically initialized *Q* at $-10q_m$ to enrich genetic diversity and hasten convergence⁵⁶. The positive-definite ρ constraint contributed a term $\sum_i n_i^2 \Theta(-n_i)$, where n_i is the number of electrons in voxel *i*. This term is zero except where the number in a voxel is negative (that is, the charge is positive). Finally, the COTR intensity constraint contributed a term $\sum_\lambda (\Delta_\lambda/\epsilon_\lambda)^2$ to the cost function, where Δ_λ is the variation in the overall COTR intensity during a reconstruction relative to the nominal measured value at each λ , and ϵ_λ is the relative uncertainty of the measurement. Weights assigned to the various terms were strong, but finite and balanced. Underweighting led to unphysical solutions, overweighting to premature termination of a reconstruction in local minima and unbalanced weighting to over-prioritization of one term at the expense of others. Nevertheless, these criteria still left a wide range of flexibility in choosing workable weighting factors. Supplementary Fig. 2 diagrammatically depicts the digital reconstruction and cost analysis procedures.

Cluster analysis. Subjective judgements of the uniqueness of reconstructed $\{\rho_{ij}\}$ for a given COTR dataset can be simply based on the visual inspection of similarities within and between small ($N \leq 10$) groups of *N* randomly initialized reconstructions. Supplementary Fig. 4 shows an example of such a group. As *N* increases, the systematic assessments of uniqueness become necessary. Here quantitative judgements are based on cluster analysis⁴⁴. Specifically, we used a *K*-means clustering algorithm⁵⁹, which partitions *N* reconstructions $\{\rho_{ij}\}$ (where $1 \leq j \leq N$) for a COTR dataset into a small fixed number $K \ll N$ of clusters in which each $\{\rho_{ij}\}$ belongs to the cluster with the nearest mean $\{\bar{\rho}_{ik}\}$, where $\bar{\rho}_{ik} = N_\kappa^{-1} \sum_{j=1}^{N_\kappa} \rho_{ij}$, N_κ is the number of reconstructions in the cluster and $\sum_{\kappa=1}^K N_\kappa = N$. A simple measure of the closeness of reconstruction *j* to $\{\bar{\rho}_{ik}\}$ is the variance $\delta_{jk} = N_v^{-1} \sum_{i=1}^{N_v} \|\rho_{ij} - \bar{\rho}_{ik}\|^2$. The average variance within cluster κ is then $\Delta_\kappa = N_\kappa^{-1} \sum_{j=1}^{N_\kappa} \delta_{jk}$. The algorithm then reassigns

reconstructions among the K clusters until global variance $\Lambda = \sum_{k=1}^K \Delta_k$ is minimized.

To visualize clusters within a set of N reconstructions, we generated two-dimensional plots (for example, Fig. 4e) using PCA⁶⁰. In PCA, each $\{\rho_{ij}\}$ is treated as a point in N_v -dimensional space. Principal components are unit vectors, where the m th vector is the direction of a line that best fits the N points and being orthogonal to the first $m - 1$ vectors. Here best fit means that the line minimizes the average squared perpendicular distance of the N points to the line. The first principal component is, thus, a line along the direction of maximum variance of the N points. The second principal component defines the direction of maximum variance in what is left once the effect of the first component is removed. Figure 4e plots the reconstruction results with respect to the first two principal components only. Such plots help us to visualize the clusters of closely related points. This identification then provides a basis for evaluating the difference between solutions within each cluster and between different clusters.

In a K -means cluster analysis⁵⁹, K is an input parameter, that is, the algorithm identifies the designated number of clusters, assigning each point $\{\rho_{ij}\}$ in a way that minimizes the global variance Λ . Designating $K = 2$ usually generated two well-separated clusters on a PCA plot of reconstruction sets $\{\rho_{ij}\}$. Members of these two clusters were approximately longitudinal mirror images of each other, that is, one is nearly a reflection of the other from a plane perpendicular to the z axis. Supplementary Figs. 3d and 4 show examples of this. This happens because an e^- bunch and its longitudinally inverted counterpart produce identical multispectral COTR data. Thus, any set of N reconstructions divides naturally and unavoidably into these two indistinguishable clusters of solutions. In principle, phase-sensitive measurements could distinguish them. However, the intensity measurements used here can, at best, unfold the e^- -beam structure to within a longitudinal mirror image of itself.

For occasional down-ramp-injected shots, a third well-separated cluster emerged and lowered Λ when we designated $K = 3$. Figure 4 is an example of this. In most such cases, the elements of the third cluster retained the main structural characteristics of the other clusters and was related to them through a simple geometric transformation. Supplementary Fig. 4 shows the details. Rarely, however, did further insights emerge from designating $K \geq 4$. The main indicators that the maximum useful K had been designated for a given dataset were (1) designating a higher K did not yield a well-separated additional cluster on a PCA plot; (2) designating a higher K did not lower Λ ; (3) cluster-averaged structures fit the COTR data as well as (or nearly as well as) individual reconstructions; and (4) visual inspection revealed greater structural differences between clusters than within a cluster.

Indicator (3) above provided our formal metric for uniqueness. Most down-ramp-injected shots yielded two or three clusters of reconstructions in each of which NSS of the cluster-averaged structure was nearly as small as that of individual reconstructions. This indicated that cluster members possessed common features that survived averaging. We relied on visual inspection to evaluate the geometric relationship between clusters. When, as in most cases, this relationship was simple, these reconstructions were deemed unique. Most self-injected shots, on the other hand, yielded clusters in which the NSS of the cluster-averaged structure substantially exceeded that of individual reconstructions, indicating that the reconstruction had not found a unique solution. Supplementary Figs. 7 and 8 show representative examples of this.

Data availability

Experimental data were generated at HZDR's DRACO facility. A collection of unprocessed experimental data is available via HZDR's ROssendorf DAta REpository (RODARE) at <https://doi.org/10.14278/rodare.2856> (ref. 61). Source data for Figs. 1–5 and Supplementary

Figs. 1–8 are available via RODARE at <https://doi.org/10.14278/rodare.2991> (ref. 62). Additional inquiries about the data should be directed to M.L. or the corresponding author.

Code availability

The relativistic PIC code PIConGPU, which supports the results of this study, is open source and freely available via GitHub. It is developed and maintained by the Institute for Radiation Physics at HZDR in close collaboration with the Center for Advanced Systems Understanding (CASUS). PIConGPU is fully documented within this Article, the Supplementary Information and ref. 55. Additional inquiries about the codes should be directed to A.D. (a.debus@hzdr.de), R.P. (r.pausch@hzdr.de) and J.T. (j.tiebel@hzdr.de) The differential evolution code used to reconstruct the 3D charge density is available via RODARE at <https://doi.org/10.14278/rodare.2856> (ref. 61). Additional inquiries about this code should be directed to M.L. (max.laberge@utexas.de, m.laberge@hzdr.de).

References

- Schramm, U. et al. First results with the novel petawatt laser acceleration facility in Dresden. *J. Phys.: Conf. Ser.* **874**, 012028 (2017).
- Schöbel, S. et al. Effect of driver charge on wakefield characteristics in a plasma accelerator probed by femtosecond shadowgraphy. *New J. Phys.* **24**, 083034 (2022).
- Bussmann, M. et al. Radiative signatures of the relativistic Kelvin-Helmholtz instability. In *Proc. International Conference on High Performance Computing, Networking, Storage and Analysis 5* (ACM, 2013).
- Lehe, R., Lifschitz, A., Thaur, C., Malka, V. & Davoine, X. Numerical growth of emittance in simulations of laser-wakefield acceleration. *Phys. Rev. ST Accel. Beams* **16**, 021301 (2013).
- Boris, J. Relativistic plasma simulation—optimization of a hybrid code. In *Proc. Fourth Conference on Numerical Simulation of Plasmas 3* (1970).
- Esirkepov, T. Z. Exact charge conservation scheme for particle-in-cell simulation with an arbitrary form-factor. *Comput. Phys. Commun.* **135**, 144–153 (2001).
- Pausch, R. & Chang, Y.-Y. Simulation code PIConGPU and setup for 'Reduction of the electron beam divergence of laser wakefield accelerators by integrated plasma lenses'. *RODARE* <https://rodare.hzdr.de/record/2361> (2023).
- Storn, R. On the usage of differential evolution for function optimization. In *Proc. North American Fuzzy Information Processing 519–523* (1996).
- Castellano, M. & Verzilov, V. Spatial resolution in optical transition radiation beam diagnostics. *Phys. Rev. ST Accel. Beams* **1**, 062801 (1998).
- LaBerge, M. *Coherent optical diagnostics of laser-wakefield-accelerated electron bunches*. PhD dissertation, Univ. of Texas–Austin (2022).
- Kanungo, T. et al. An efficient K -means clustering algorithm: analysis and implementation. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**, 881–892 (2002).
- Jolliffe, I. & Cadima, J. Principal component analysis: a review and recent developments. *Phil. Trans. R. Soc. A* **374**, 20150202 (2016).
- LaBerge, M. et al. Data publication: Revealing the 3D structure of microbunched plasma-wakefield-accelerated electron beams. *RODARE* <https://doi.org/10.14278/rodare.2856> (2024).
- LaBerge, M. et al. Source data: Revealing the 3D structure of microbunched plasma-wakefield-accelerated electron beams. *RODARE* <https://doi.org/10.14278/rodare.2991> (2024).

Acknowledgements

M.L., B.B., A.H., R.Z. and M.C.D. acknowledge support from the US Department of Energy grants DE-SC0011617 (M.C.D.) and DE-SC0014043 (M.C.D.) and the US National Science Foundation

grant PHY-2308921 (M.C.D.). M.C.D. acknowledges additional support from the Alexander von Humboldt Foundation. A.H.L. acknowledges support from the Fermi Research Alliance under contract no. DE-AC02-07CH11359 with the US Department of Energy. M.L., Y.-Y.C., J.C.C., A.D., R.P., S.S., J.T., P.U., A.W., O.Z., U.S. and A.I. acknowledge support from the Helmholtz Association under the Accelerator Research and Development (ARD) topic of the Helmholtz Matter and Technologies program (U.S.).

Author contributions

M.L. designed and constructed the multispectral COTR apparatus, led the acquisition and analysis of experimental data and drafted the manuscript. B.B., Y.-Y.C., J.C.-C., A.H., S.S., P.U., R.Z. and O.Z. assisted with the data acquisition. A.I. directed the experimental activity in coordination with the DRACO facility staff and supervised the student participants. R.P., A.D. and J.T. carried out the PIC simulations and coordinated with M.L., A.I., M.C.D., R.Z. and A.H.L. on interpreting the experimental results. A.W. led the statistical cluster analysis of the reconstructed e^- bunch solutions. M.C.D. conceived the experiments, inspired by past COTR research by A.D., O.Z., A.I., U.S. and A.H.L., supervised the work of M.L., B.B., A.H. and R.Z., and wrote the final

version of the paper in consultation with M.L., A.I., R.P., A.D., U.S. and A.H.L. All authors discussed the results and commented on the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41566-024-01475-2>.

Correspondence and requests for materials should be addressed to M. C. Downer.

Peer review information *Nature Photonics* thanks David Attwood, Makina Yabashi and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.