Single-shot, transverse self-wakefield reconstruction from screen images

N. Majernik^(b),^{1,*,†} W. Lynn,¹ G. Andonian^(b),¹ T. Xu^(b),^{2,†} P. Piot^(b),^{2,3} and J. B. Rosenzweig¹

¹University of California Los Angeles, Los Angeles, California 90095, USA ²Northern Illinois University, DeKalb, Illinois 60115, USA ³Argonne National Laboratory, Lemont, Illinois 60439, USA

(Received 12 June 2023; accepted 5 September 2023; published 25 September 2023)

A single-shot method to reconstruct the transverse self-wakefields acting on a beam, based only on screen images, is introduced. By employing numerical optimization with certain approximations, a relatively high-dimensional parameter space is efficiently explored to determine the multipole components of the transverse-wakefield topology up to desired order. The reconstruction technique complements simulations, which are able to directly describe the wakefield composition based on experimental conditions. The technique is applied to representative simulation results as a benchmark, and also to experimental data on wakefield observations driven in dielectric-lined structures.

DOI: 10.1103/PhysRevAccelBeams.26.094601

I. INTRODUCTION

Electromagnetic wakefields generated by relativistic particle beams propagating in slow-wave structures or near resistive boundaries are an important consideration in many aspects of accelerator and beam physics research. The transverse components of self-generated wakefields that act back on the beam may be desirable and intentionally excited, e.g., as passive streakers for beam diagnostics [1] or in emerging wakefield-based focusing schemes [2,3]. However, transverse self-wakefields may also lead to emittance growth and possibly seed a beam-breakup instability [4,5]. Detailed information about the wakefields and the subsequent beam dynamics are thus critical for accelerator design and, in particular, advanced acceleration concepts based on wakefields in plasma or structures. Due to the spatial and temporal scales involved, as well as the relativistic nature of the beam and its induced fields, it is challenging to directly interrogate self-wakefields in situ. Nevertheless, it is possible to ascertain properties of the electromagnetic wakefields based on the measured frequency content of the outcoupled radiation [5,6], or by comparing experimental measurements of beam properties to simulations. However, a method to directly determine the detailed characteristics of wakefields themselves is much more desirable. In a plasma wakefield accelerator, it may be possible to directly image the plasma wave and thus infer information about the wakefields either optically [7–11] or using a secondary electron beam [12]. Such nondestructive approaches are not applicable to wake interactions which occur in vacuum structures within material enclosures (e.g., wakefields from dielectric lined or corrugated metallic structures). Other wakefield-based techniques have been proposed which may be more appropriate for in-vacuum interactions, including using tailored probe beams [13]. In this paper, we introduce a new technique for investigating the transverse self-wakefields without requiring any additional experimental infrastructure, by measuring, on a standard screen, the two-dimensional transverse projections [in the (x, y) plane] of the beam undergoing transverse self-wake interactions.

In order to develop a foundation for the reconstruction approach, we revisit the formalism of wakefield descriptions using a standard multipole decomposition. It is possible to describe the transverse wakefields, W, expressed in terms of the net transverse momentum kick, at each longitudinal position along the beam with a potential-based formalism [14]:

$$A + iV = \sum_{n=1}^{\infty} (a_n + ib_n)(x + iy)^n,$$
$$W = -\left\{\frac{\partial V}{\partial x}, \frac{\partial A}{\partial x}\right\}.$$
(1)

This yields transverse wakefields of the form

$$W = -\left\{\frac{\partial V}{\partial x}, \frac{\partial A}{\partial x}\right\}$$

= $\{-b_1 - 2a_2y - 2b_2x - 6a_3xy + 3b_3(-x^2 + y^2) + \cdots, -a_1 - 2a_2x + 2b_2y + 3a_3(-x^2 + y^2) + 6b_3xy + \cdots\},$
(2)

majernik@slac.stanford.edu

[†]Present address: SLAC National Accelerator Laboratory, Menlo Park, California 94025, USA.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

where a_n and b_n are the *n*th-order skew and normal multipole coefficients, respectively. By confining the reconstruction to low order multipole moments (an approximation that is quantified and justified below) it is possible to infer the self-wakefields by examining the change that they cause in the transverse distribution of a beam imaged on a screen.

Furthermore, in this work we utilize a thin-lens approximation, i.e., that the beam spatial distribution does not change significantly over the length of the wakefield interaction itself, and thus only the integrated kick from the interaction (in the instantaneous approximation) is reconstructed.

We will also assume that the transverse wakefield at every longitudinal position along the beam (denoted simply as z, with beamline distances indicated by s) is proportional up to a multiplicative constant, $\kappa(z)$, which is a normalizing factor determined from the data. The approximation holds provided that the bunch length is short compared to the wavelength of all modes which appreciably contribute to the overall wakefield and will be further justified below. In this limit, each mode's convolution with the beam, over the duration of the beam, is well approximated as a curve which is monotonically increasing toward the tail of the beam. Though each of these contributions will have different slopes depending on their wave number, since they all start at zero (by construction) their ratios will be constant, justifying this approximation.

In this paper, we first describe the mathematical algorithm to reconstruct transverse self-wakefields. We then use the algorithm to reconstruct wakefields using simulation results as a benchmark. The technique is then extended to experimental data from a recent dielectric wakefield study conducted at the Argonne Wakefield Accelerator (AWA) facility. Finally, we discuss the range of applicability for the reconstruction and some simple ways to extend the technique's fidelity.

II. TRANSVERSE WAKEFIELD RECONSTRUCTION

The goal of the reconstruction method (Fig. 1) is to fully describe the transverse self-wakefields affecting the beam. The input for the technique is two images, taken on the same screen downstream of the wakefield interaction point: one with the structure in place, and one without. Herein, "structure" is used generically and can refer to either a plasma, dielectric-lined waveguide, corrugated waveguide, or any other wakefield-generating element that acts back on the beam. The screen images are used as probability density functions (PDF) which are sampled to create two populations of n, (x, y) macroparticles. A 2D grid of bins is defined which is large enough to cover both beam distributions, with a bin spacing small enough to resolve the beams' structures. The wakefield-affected (i.e., "withstructure") macroparticles are binned, yielding bin-wise counts $B_{i,j}$. The without-structure macroparticles are then



FIG. 1. Schematic overview of the algorithm for the single-shot reconstruction of transverse beam wakes from screen images. The bin size has been exaggerated for clarity.

transformed according to $(x, y) + (x_c, y_c) + \kappa \mathbf{W}L/p_z$ where (x_c, y_c) are corrections for beam drift between the without-structure and with-structure shots (omitted for simulations where this effect is not present), κ is the wakefield amplitude scale factor, sampled independently from the PDF of the random variable \mathcal{K} , L is the effective drift length from the structure to the screen, and p_z is the mean longitudinal momentum of the beam. The transformed macroparticles are binned, yielding $\tilde{B}_{i,j}$. We define the overall error as the sum of the absolute errors for all the bins, divided by twice the number of macroparticles to avoid double counting:

$$\epsilon_{\rm fit} = \sum_{i,j} \frac{|\tilde{B}_{i,j} - B_{i,j}|}{2n},\tag{3}$$

which ranges from zero to one. A numerical optimizer is employed to minimize this error term by finding the best parameter vector $\{x_c, y_c, \mathcal{K}_1, \mathcal{K}_2, ..., a_1, b_1, a_2, b_2, ...\}$ where $\mathcal{K}_1, \mathcal{K}_2, ...$ parameterize \mathcal{K} , with multipole coefficients a_i , b_i up to the selected order. The technique is visually summarized in Fig. 1.

The random variable \mathcal{K} results from weighting the probability of the wakefield strength at a longitudinal position, $\kappa(z)$, by the current at that position, I(z). Depending on the available information and the expected complexity, different definitions for the PDF may be used. For the simulated cases below, the values of $\kappa(z)$ and I(z) can be found directly and used to construct the PDF. For experiments for which there are trusted simulations, this approach may also be used. However, it is also possible to use less prior knowledge. Shown is an example of a convenient and versatile PDF, used by several reconstructions below, with one optimizer parameter, \mathcal{K}_1 :

$$f_{\mathcal{K}}(\kappa) = \begin{cases} (1 + \mathcal{K}_1)\kappa^{\mathcal{K}_1} & 0 < \kappa < 1\\ 0 & \text{otherwise} \end{cases}.$$
 (4)

If this does not give convergent results, it is possible to define the distribution based on other assumptions or with a greater number of free parameters.

Several methods were attempted for optimization in this relatively high dimensional parameter space. Derivativebased optimizers encountered difficulties due to the discontinuous and noisy nature of the data, resulting from binning discrete particles. Bayesian optimization methods performed better, but the comparatively high computational overhead of this technique made it suboptimal. The best performance was found using a derivative-free [15] global optimization, in particular, differential evolution. This approach is not guaranteed to arrive at the complete wakefields, but since the relevant wakefields are often dominated by a few, low degree multipole components, it serves as an efficient and effective approximation. Alternative options may include, instead of binning the particles to discrete bins, using a differentiable optimization based on kernel density estimation [16].

III. BENCHMARKING TO SIMULATION

A variety of dielectric wakefield scenarios were simulated using the finite-difference time-domain (FDTD) particle-in-cell (PIC) code WARP [17] (summarized in the Appendix, in Table I). These scenarios are inspired by recent experimental work which investigated wakefields arising from a variety of beam distributions, including those containing skew components.

In the first scenario chosen, a beam with a charge of 1.95 nC, bunch size $\sigma_x = 1.76$ mm, $\sigma_y = 0.23$ mm, $\sigma_z = 0.64$ mm, and a tilt of 2.2° (see Fig. 2) was propagated past an alumina slab ($\epsilon_r = 9.9$) 5 mm thick and 150 mm long, coated on the side opposite the beam with a layer of metal. The beam centroid was 0.91 mm below the surface of the dielectric. This simulation is most similar to the recent wakefield experiment performed at AWA [18] with the exception that the simulated beam energy was 500 MeV. This permitted us to freeze the motion of the beam during the interaction to ensure the validity of the thin lens assumption. The transverse momentum of each particle was recorded before and after the simulated interaction and used to construct an interpolated function for the transverse wakefield $W_{sim}(x, y, z)$, taken to be the ground truth in the analysis.

 $W_{sim}(x, y, z)$ is then sampled at a number of z slices, spaced with $\Delta z \ll \sigma_z$. These fields were fit with the best W_{guess} by finding the vector $c_{slice} = \{a_1, b_1, a_2, b_2, ..., a_n, b_n\}$ which minimizes

$$\langle |\boldsymbol{W}_{\text{sim}}(\boldsymbol{x}_k, \boldsymbol{y}_k, \boldsymbol{z}_{\text{slice}}) - \boldsymbol{W}_{\text{guess}}(\boldsymbol{x}_k, \boldsymbol{y}_k)| \rangle_{95\%}, \qquad (5)$$

where *k* counts along the simulation macroparticles and $\langle \cdot \rangle_{95\%}$ is a truncated mean which excludes 5% of the population with the worst errors to prevent stray particles from unduly affecting the result. In this case, coefficients up to n = 6 (dodecapole) were used. Using these fit coefficients for each slice and then linearly interpolating between the slices yields $W_{\text{slicewise}}$. As a confirmation that a high enough multipole order was appropriate for the fit, the slicewise coefficients are compared back to the ground truth wakefield:

$$\epsilon_{\text{slicewise}} = \frac{\langle |\boldsymbol{W}_{\text{sim}}(\boldsymbol{x}_k, \boldsymbol{y}_k, \boldsymbol{z}_k) - \boldsymbol{W}_{\text{slicewise}}(\boldsymbol{x}_k, \boldsymbol{y}_k, \boldsymbol{z}_k)| \rangle_{95\%}}{\langle |\boldsymbol{W}_{\text{sim}}(\boldsymbol{x}_k, \boldsymbol{y}_k, \boldsymbol{z}_k)| \rangle_{95\%}}.$$
(6)

Equation (6) yields a relative error of 1.5% indicating that, for this case, n = 6 is sufficient to describe the transverse wakefields. $\epsilon_{\text{slicewise}}$ will improve monotonically with greater *n* but with diminishing returns.



FIG. 2. Figures relating to the 500 MeV, top slab, with PDF WARP simulation and reconstruction (see Table I) (a) Transverse distribution of without-structure beam. (Plot range is constant for all subfigures. Charge density (arb. units) color scale bar is constant between distributions). (b) Reconstructed wakefields, W_{opt} . White region contains 90% of beam charge. (c) Implied with-structure beam distribution by W_{opt} . (d) Ground truth wakefields, W_{agg} . (e) Ground truth with-structure beam distribution.

In order to find the *aggregate multipole* coefficients, c_{agg} , calculate the average of all the c_{slice} vectors, weighted by the current at that longitudinal position, $I(z_{slice})$

$$\boldsymbol{c}_{\text{agg}} = \sum_{\text{slices}} I(\boldsymbol{z}_{\text{slice}}) \boldsymbol{c}_{\text{slice}} / \sum_{\text{slices}} I(\boldsymbol{z}_{\text{slice}}). \tag{7}$$

These coefficients, via Eq. (1), yield the aggregate wake-field W_{agg} .

As previously discussed, we assume that the fields at every slice are equal to the aggregate wakefield, including a multiplicative constant: $c_{\text{slice}} \approx \kappa(z) c_{\text{agg}}$. The approximation is valid if the wavelengths of modes which produce relevant multipoles are long relative to σ_z , which is reasonable for the scenarios under consideration. In Fig. 3, the normalized, slicewise coefficients for the first three multipole moments are shown as a function of z. Since the lines for the dipole, quadrupole, and sextupole orders are approximately the same shape within the area of most of the beam charge, this is evidence that approximating the fields as constant, up to a z-dependent multiplicative factor, is reasonable. $\kappa(z)$ is calculated by finding the value which minimizes the difference between $\kappa(z_{\text{slicewise}})W_{\text{agg}}$ and $W_{\text{slicewise}}$ at each z_{slice} and linearly interpolating between these values.

As a final check of both sufficient multipole order and the $c_{\text{slice}} \approx \kappa(z)c_{\text{agg}}$ assumption, $\kappa(z)c_{\text{agg}}$ is compared to the ground truth at all the macroparticle positions:

$$\varepsilon_{\text{agg}} = \frac{\langle |\boldsymbol{W}_{\text{sim}}(\boldsymbol{x}_k, \boldsymbol{y}_k, \boldsymbol{z}_k) - \kappa(\boldsymbol{z}_k) \boldsymbol{W}_{\text{agg}}(\boldsymbol{x}_k, \boldsymbol{y}_k)| \rangle_{95\%}}{\langle |\boldsymbol{W}_{\text{sim}}(\boldsymbol{x}_k, \boldsymbol{y}_k, \boldsymbol{z}_k)| \rangle_{95\%}}.$$
 (8)

Equation (8) yields an overall relative error of 2.3%. Note that $\epsilon_{\text{slicewise}}$ (1.5% in this case) will set a lower bound on



FIG. 3. At each longitudinal position along the beam, a multipole decomposition is performed. The values of $\sqrt{a_i^2 + b_i^2}$ are shown, after being normalized. Also shown is the current profile of the beam, with the white region near the center corresponding to 90% of the beam charge.



FIG. 4. A histogram showing the probability density function of κ , the multiplicative weighting factor. These results are for the 500 MeV, top slab WARP simulation with the "Exact PDF" corresponding to using the scaled I(z) weighted exact $\kappa(z)$ and the "Best fit PDF" corresponding to Eq. (4) with $\mathcal{K}_1 = -0.35$, determined using the reconstruction algorithm.

 ϵ_{agg} ; the close agreement between these values indicate that the longitudinal approximation is reasonable. Specifically, we have managed to capture the overwhelming majority of the information regarding the 3D map of the transverse wakes (**W**: $\mathbb{R}^3 \to \mathbb{R}^2$) using only 12 coefficients for the multipoles and a 1D function, $\kappa(z)$. The PDF for the wakefield strength was constructed from $\kappa(z)$ and I(z); it is shown in Fig. 4.

Simulated screen images, 3 m downstream of the exit of the structure, were produced from the simulation data and fed into the optimizer described in the previous section. Initially, the kick PDF was provided so the only free parameters were the multipole coefficients which were permitted up to n = 6. Figure 2 shows the without-structure beam and reconstructed with-structure beam and associated fields, as well as the ground truth with-structure beam and fields. The optimizer figure of merit, $e_{\rm fit}$, converged to 1.3%. Finally, the ground truth aggregate wakefield, $W_{\rm agg}$, is compared to the reconstructed wakefield, $W_{\rm opt}$:

$$\epsilon_{\text{recons}} = \left\langle \frac{|\boldsymbol{W}_{\text{agg}}(\boldsymbol{x}_{k}, \boldsymbol{y}_{k}) - \boldsymbol{W}_{\text{opt}}(\boldsymbol{x}_{k}, \boldsymbol{y}_{k})|}{\langle |\boldsymbol{W}_{\text{agg}}(\boldsymbol{x}_{k}, \boldsymbol{y}_{k})| \rangle_{95\%}} \right\rangle_{\text{median}}, \quad (9)$$

where $\langle \cdot \rangle_{\text{median}}$ is the median. For this scenario, the reconstruction technique described here recovered the ground truth wakefield to within 2.7%.

This particular reconstruction assumes that the PDF of \mathcal{K} is known. If there are accurate simulations available for the experiment under consideration, this is of course possible. More generally though, we anticipate that the \mathcal{K} will not be known. We repeat the reconstruction of this scenario with one more free parameter, \mathcal{K}_1 , and the PDF from Eq. (4).



FIG. 5. Reconstruction of 500 MeV, top slab WARP simulation under different reconstruction conditions. Insets show the wakefields that result in the corresponding beam. For additional details refer to Table I. (a) Without-structure beam. (b) Not given κ PDF, reconstruction to n = 2. (c) Not given κ PDF, reconstruction to n = 4. (d) Not given κ PDF, reconstruction to n = 6. (e) Given κ PDF, reconstruction to n = 6. (f) Ground truth, with-structure beam.

The final result is quite similar; the reconstructed PDF is scaled and shown compared to the ground truth PDF in Fig. 4 and the with-structure beam is shown in Fig. 5. Before calculating ϵ_{recons} , the same multiplicative scaling factor that was used to adjust the PDF is applied to W_{opt} ; this does not change the reconstructed fields. By approximating the PDF with Eq. (4) the overall reconstruction error, ϵ_{recons} , is degraded only slightly, to 3.7%.

Since the ground truth to which the reconstructions are compared has been defined to finite order, by construction, there is no advantage to giving the reconstruction algorithm additional, free multipole coefficients beyond n = 6. However, for experimental data, it will not be known a priori to what order reconstruction should be performed. Instead, the number of free parameters should be chosen such that there are diminishing returns for a higher order. This process has been tested on simulated data and summarized in Fig. 6. Here it can be observed that both optimizer figure of merit ϵ_{fit} and the overall reconstruction error $\epsilon_{\rm recons}$ trend down with increasing reconstruction order but reach diminishing returns around n = 4 or 5. The effect of changing the conditions of the reconstruction, including varying n can be visually inspected in Fig. 5. This same approach will be used for experimental data; recall that $\epsilon_{\rm fit}$ will be available for reconstructions of experimental data but $\epsilon_{\rm recons}$ cannot be known. Using a higher number of free parameters will dramatically increase the computational expense and also increase the possibility of overfitting the data, leading to worse estimates.

Further simulations were conducted using the same beam parameters, but reducing the energy from 500 to 42 MeV to more closely align with the experimental scenario at AWA. This change will reduce the validity of the thin lens approximation, as the less rigid particles move during the wakefield interaction, causing both the particle distribution and wake distribution to evolve along



FIG. 6. The effect of changing the reconstruction order, n, on the error metrics $\epsilon_{\rm fit}$ and $\epsilon_{\rm recons}$ for the 500 MeV, top slab WARP simulation. Diminishing returns for both metrics are achieved around n = 4 or 5.

the structure. In this case, the virtual screen is 0.2 m downstream of the end of the structure. But, since we now must assume that the beam evolves within the structure the effective drift, L, also includes one half of the structure 15 cm structure length, i.e., L = 0.28 m. We still assume that the evolution within the structure is sufficiently small that the wakefields do not change. As expected, the lower energy beam does reduce the accuracy of the model (summarized in Table I). Crucially though, since the difference from the ground truth, $\epsilon_{\rm recons}$, is only 4.3% at 42 MeV, this is still a very high-quality reconstruction and illustrates the flexibility of this approach and its applicability to experimental scenarios of interest. In a later section, we will discuss some possible extensions to the model which could improve its reconstructions further from the thin lens limit.

This same 42 MeV simulated beam is shown reconstructed in several ways in Fig. 7. In the first case, it is



FIG. 7. Reconstruction of without-structure beam, (a), into different with-structure beams. This illustrates how the technique is single-shot, once a without-structure image has been collected. The first row is the without-structure beam, the second row is the ground truth, with-structure beam, the third row is the reconstructed wakefields, and the final row is the reconstructed with-structure beams. Both cases start from the same without-structure beam, (a). The left column, (b–d), corresponds to the 42 MeV, top slab, without PDF simulation while the right column, (e–g), corresponds to the 42 MeV, double slab, without PDF simulation. The insets on (b) and (e) show the ground truth wakefields from each simulation (scale bars = 1 mm).

propagated past a single dielectric slab on the top, while in the second case, it passes on the symmetry plane between a pair of slabs: one on top and one on the bottom. These reconstructions were accomplished using the same withoutstructure screen image, illustrating the single-shot nature of the technique: once a without-structure image is acquired, provided the beam is stable over time, all subsequent withstructure screen images can be reconstructed on a singleshot basis.

IV. RECONSTRUCTING EXPERIMENTAL DATA

This technique has also been applied to experimental data, summarized in Table I. In Fig. 8, a dielectric wakefield



FIG. 8. Reconstruction of experimental data with Top slab, 7.2° tilt beam case under different reconstruction conditions. For additional details refer to Table I. (a) Measured without-structure beam. (b) Measured with-structure beam. (c and d) Reconstructed wakefields and implied with-structure beam to n = 6. (Reconstructed $\mathcal{K}_1 = -0.70$) (e and f) Reconstructed wakefields and implied with-structure beam to n = 3. (g and h) Reconstructed wakefields and implied with-structure beam to n = 1.

interaction at the AWA between a 7.2° tilted, 42 MeV beam has been reconstructed up to sixth multipole order with lower order reconstructions also shown. This effort is discussed in much greater detail in [18] where the reconstructed wakefields are compared to both simulations and analytic models of the skew wake interaction. In this case, the PDF of \mathcal{K} was estimated with Eq. (4). x_c and y_c were also varied, for a total parameter space dimension of 15. The final reconstruction agrees with the actual withstructure beam image to within 6.0%, suggesting a successful inference of the true wakefields.

Naturally, shot-to-shot and long term drifts in the beamline will reduce the quality of the reconstruction. Obviously, if the without-structure reference beam used in the reconstruction algorithm does not remain substantially similar to the actual beam which entered the structure, the reconstruction will be less accurate. The inclusion of free parameters to remove jitter, x_c and y_c , is able to address only the most basic sort of jitter, and not longer term drifts.

V. DISCUSSION

We have introduced and explained a technique for reconstructing the transverse self-wakefields of a beam without additional experimental infrastructure. The approach has been benchmarked against simulations where the simulated fields may serve as a ground truth, enabling us to quantify the accuracy of the technique. We have, in this way, found agreement at the few-percent level for experimentally relevant scenarios. This technique, with its underlying methods explained in detail, has also been used to reconstruct the transverse self-wakes for experimental data. Comparison to simulation and theory for the particular experiment in question is given elsewhere [18].

With improvements to the implementation of the technique, specifically in the numerical optimizer, it may be possible to use this technique as a real-time diagnostic. By combining the reconstructed wakefield information with advanced beam control (e.g., multileaf collimators [19] or a range of other options [20]) possibly also leveraging additional non-destructive diagnostic information [21–23] it should be possible to exert very precise control over the wakes to establish a desired condition.

Although this technique has been demonstrated to be useful for a range of experimentally relevant scenarios, future extensions may further increase its range of applicability. One such step that would further loosen the thin lens requirement would be to reconstruct in multiple steps, evolving the particle distribution (such that each particle now can sample different parts of the wakefield) along the structure but assuming that the wakefield functions themselves are unchanging. It may also be possible to fit multiple κ PDFs for different moments to loosen the requirement that the beam be much shorter than all relevant mode wavelengths. For a bunch experiencing energy changes from longitudinal wakefields but otherwise continuing to satisfy the thin-lens approximation, the described method is still valid if knowledge of the absolute wakefield magnitude is not required. Our approach of describing the wakefields using a PDF for the strength is adequate without further change, provided the energy change is constant within a *z*-slice, since the effect of reduced energy is indistinguishable from increased wakefield strength. As such, the effect of this change in energy would continue to be encapsulated by the strength PDF. With additional information, e.g., a spectrometer, the absolute magnitude of the wakes could be recovered.

In a wakefield accelerator, any accelerated bunch will necessarily be far enough from the drive bunch to violate the assumption that the total current profile is short compared to the structure modes. However, it may be possible to reconstruct the wakes for the drive bunch and witness bunch independently. Specifically, this may be approached by running a drive bunch alone with and without the structure present; this technique can be applied directly to reconstruct its self-wakes. Then, the drive and witness can be run through the structure and the drive beam's contribution to the screen images may be background subtracted in an effort to capture only the witness beam's with-structure and without-structure distributions. Subsequently, this reconstruction algorithm may be applied to the background-subtracted, witness-only images. This would be exciting future work and could function synergistically with patterned witness beams, e.g., [13]. By using short witness beams and adjusting the delay relative to the drive beam, it would also be possible to measure the transverse wakefields at different longitudinal positions to reconstruct the full 3D wakefields.

ACKNOWLEDGMENTS

We gratefully acknowledge the valuable support from the staff of the Argonne Wakefield Accelerator and fruitful discussions with S. Baturin and A. Halavanau. This work was supported by DOE Grants No. DE-SC0017648 and No. DE-SC0009914. T. X. and P. P. are supported by DOE Grant No. DE-SC0022010.

APPENDIX: SUMMARY OF RECONSTRUCTIONS

Reconstruction details and figures of merit for selected cases are summarized in Table I.

TABLE I. Summary of figures of merit for reconstructions of simulated and experimental screen images. $\epsilon_{\text{slicewise}}$ [Eq. (6)] compares the slicewise multipole fits to the ground truth wakes. ϵ_{agg} [Eq. (8)] compares the aggregate multipole coefficients and longitudinally dependent multiplicative factor, $\kappa(z)$, to the ground truth wakes. ϵ_{fit} [Eq. (3)] is the figure of merit for the optimizer and expresses the error between the binned without-structure beam, transformed according to the best parameter vector, and the with-structure beam. ϵ_{recons} [Eq. (9)] compares the reconstructed wakes to the ground truth wakes.

Name	Data type	$\epsilon_{ m slicewise}$ (%)	$\epsilon_{ m agg}$ (%)	Multipole Total free order, <i>n</i> parameters ϵ_{fit} (%)			$\epsilon_{\rm recons}$ (%)
500 MeV, top slab, with PDF				6	12	1.3	2.7
500 MeV, top slab, without PDF	Simulation	1.5	2.3	6	13	1.8	3.7
				5	11	2.1	4.0
				4	9	4.0	9.9
				3	7	7.1	10
				2	5	14	24
				1	3	20	65
				0	0	30	100
42 MeV, top slab, with PDF		2.4	3.2	6	12	1.5	4.3
42 MeV, top slab, without PDF				6	13	1.7	4.6
42 MeV, double slab, without PDF		2.2	2.6	6	13	1.7	4.1

(Table continued)

Name	Data type	$\epsilon_{ m slicewise}$ (%)	$\epsilon_{ m agg}~(\%)$	Multipole Total free $_{agg}$ (%) order, <i>n</i> parameters ϵ_{fit} (%)					
Top slab, 7.2° tilt	_	-		6 15 6.0	6.0				
		-			13	6.0	-		
		-		4	11	6.2	-		
				3	9	6.2			
	Experiment	Experiment Not applicable Not applicable			7	6.4	Not applicable		
	*				5	15			
					2	25			
Bottom slab, 2.2° tilt					15	8.9	_		
Top slab, 2.2° tilt				6	15	6.1	-		

TABLE I. (Continued)

- S. Bettoni, P. Craievich, A. Lutman, and M. Pedrozzi, Temporal profile measurements of relativistic electron bunch based on wakefield generation, Phys. Rev. Accel. Beams 19, 021304 (2016).
- [2] W. Lynn *et al.*, Demonstration of transverse stability in an alternating symmetry planar dielectric structure, in *Proceedings of the 14th International Particle Accelerator Conference (IPAC'23)* (JACoW Publishing, Geneva, Switzerland, 2023), pp. 972–976.
- [3] W. Lynn, G. Andonian, D. Doran, S. Kim, N. Majernik, S. OTool, P. Piot, J. Power, J. Rosenzweig, and E. Wisniewski, Transverse stability in an alternating symmetry planar dielectric wakefield structure, in *Proceedings* of the 5th International Particle Accelerator Conference (JACoW Publishing, Geneva, Switzerland, 2022), https:// doi.org/10.18429/JACoW-NAPAC2022-TUPA82.
- [4] C. Li, W. Gai, C. Jing, J. Power, C. Tang, and A. Zholents, High gradient limits due to single bunch beam breakup in a collinear dielectric wakefield accelerator, Phys. Rev. ST Accel. Beams 17, 091302 (2014).
- [5] B. D. O'Shea, G. Andonian, S. Baturin, C. I. Clarke, P. Hoang, M. J. Hogan, B. Naranjo, O. B. Williams, V. Yakimenko, and J. B. Rosenzweig, Suppression of Deflecting Forces in Planar-Symmetric Dielectric Wakefield Accelerating Structures with Elliptical Bunches, Phys. Rev. Lett. **124**, 104801 (2020).
- [6] N. Majernik, G. Andonian, O. Williams, B. O'Shea, P. Hoang, C. Clarke, M. Hogan, V. Yakimenko, and J. Rosenzweig, Positron driven high-field terahertz waves via dielectric wakefield interaction, Phys. Rev. Res. 4, 023065 (2022).
- [7] J. Marques, J. Geindre, F. Amiranoff, P. Audebert, J. Gauthier, A. Antonetti, and G. Grillon, Temporal and Spatial Measurements of the Electron Density Perturbation Produced in the Wake of an Ultrashort Laser Pulse, Phys. Rev. Lett. **76**, 3566 (1996).

- [8] C. Siders, S. Le Blanc, D. Fisher, T. Tajima, M. Downer, A. Babine, A. Stepanov, and A. Sergeev, Laser Wakefield Excitation and Measurement by Femtosecond Longitudinal Interferometry, Phys. Rev. Lett. 76, 3570 (1996).
- [9] N.H. Matlis, S. Reed, S.S. Bulanov, V. Chvykov, G. Kalintchenko, T. Matsuoka, P. Rousseau, V. Yanovsky, A. Maksimchuk, S. Kalmykov *et al.*, Snapshots of laser wakefields, Nat. Phys. 2, 749 (2006).
- [10] A. Buck, M. Nicolai, K. Schmid, C. Sears, A. Sävert, J. M. Mikhailova, F. Krausz, M. C. Kaluza, and L. Veisz, Real-time observation of laser-driven electron acceleration, Nat. Phys. 7, 543 (2011).
- [11] A. Sävert, S. Mangles, M. Schnell, E. Siminos, J. Cole, M. Leier, M. Reuter, M. Schwab, M. Möller, K. Poder *et al.*, Direct Observation of the Injection Dynamics of a Laser Wakefield Accelerator Using Few-Femtosecond Shadowgraphy, Phys. Rev. Lett. **115**, 055002 (2015).
- [12] C. Zhang, J. Hua, X. Xu, F. Li, C.-H. Pai, Y. Wan, Y. Wu, Y. Gu, W. Mori, C. Joshi,*et al.*, Capturing relativistic wakefield structures in plasmas using ultrashort high-energy electrons as a probe, Sci. Rep. 6, 1 (2016).
- [13] A. Halavanau, P. Piot, and S. S. Baturin, Single-shot transverse wakefield mapping with a hollow electron beam, arXiv:2307.02776.
- [14] J. T. Tanabe, *Iron dominated Electromagnets: Design, Fabrication, Assembly and Measurements* (World Scientific Publishing Company, Singapore, 2005).
- [15] S. Koziel and X.-S. Yang, Computational Optimization, Methods and Algorithms (Springer, New York, 2011), Vol. 356.
- [16] R. Roussel, A. Edelen, C. Mayes, D. Ratner, J. P. Gonzalez-Aguilera, S. Kim, E. Wisniewski, and J. Power, Phase space reconstruction from accelerator beam measurements using neural networks and differentiable simulations, arXiv:2209.04505.
- [17] J.-L. Vay, D. P. Grote, R. H. Cohen, and A. Friedman, Novel methods in the particle-in-cell accelerator codeframework warp, Comput. Sci. Discovery 5, 014019 (2012).

- [18] W. Lynn, T. Xu, G. Andonian, S. Doran, G. Ha, N. Majernik, P. Piot, J. Power, J. Rosenzweig, C. Whiteford, and E. Wisniewski, Observation of skewed electromagnetic wakefields in an asymmetric structure driven by flat electron bunches, arXiv:2308.09137.
- [19] N. Majernik, G. Andonian, W. Lynn, S. Kim, C. Lorch, R. Roussel, S. Doran, E. Wisniewski, C. Whiteford, P. Piot *et al.*, Beam shaping using an ultrahigh vacuum multileaf collimator and emittance exchange beamline, Phys. Rev. Accel. Beams 26, 022801 (2023).
- [20] G. Ha, K.-J. Kim, J. G. Power, Y. Sun, and P. Piot, Bunch shaping in electron linear accelerators, Rev. Mod. Phys. 94, 025006 (2022).
- [21] B. O'Shea, G. Andonian, S. Barber, K. Fitzmorris, S. Hakimi, J. Harrison, P. Hoang, M. Hogan, B. Naranjo, O. Williams *et al.*, Observation of acceleration and deceleration in gigaelectron-volt-per-metre gradient dielectric wakefield accelerators, Nat. Commun. 7, 12763 (2016).
- [22] G. Andonian, D. Stratakis, M. Babzien, S. Barber, M. Fedurin, E. Hemsing, K. Kusche, P. Muggli, B. O'Shea, X. Wei *et al.*, Dielectric Wakefield Acceleration of a Relativistic Electron Beam in a Slab-Symmetric Dielectric Lined Waveguide, Phys. Rev. Lett. **108**, 244801 (2012).
- [23] R. Roussel, G. Andonian, J. Rosenzweig, and S. Baturin, Longitudinal current profile reconstruction from a wakefield response in plasmas and structures, Phys. Rev. Accel. Beams 23, 121303 (2020).