

Lawrence Berkeley National Laboratory

LBL Publications

Title

Robust electron counting for direct electron detectors with the Back-propagation counting method

Permalink

<https://escholarship.org/uc/item/2jb1m4nm>

Journal

npj Computational Materials, 12(1)

ISSN

2057-3960

Authors

Renner, Joshua
Wright, Matthew A
Bouchard, Kristofer
[et al.](#)

Publication Date

2026-05-12

DOI

10.1038/s41524-026-02126-6

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

<https://doi.org/10.1038/s41524-026-02126-6>

Robust electron counting for direct electron detectors with the Back-propagation counting method



Joshua Renner¹, Matthew A. Wright², Kristofer Bouchard²✉, Bruce E. Cohen^{3,4}, Peter Ercius³✉, Azriel Goldschmidt⁵✉, Cassio C. S. Pedroso³, Ambarneil Saha³ & Peter Denes³

Electron microscopy (EM) is a foundational tool for directly assessing the structure of materials. Recent advances in direct electron detectors have improved signal-to noise ratios via single-electron counting. However, accurately counting electrons at high flux remains challenging. We developed a new method of electron counting for direct electron detectors, Back-Propagation Counting (BPC). BPC uses machine learning techniques designed for mathematical operations on large tensors but does not require large training datasets. In synthetic data, we show BPC is able to count multiple electron strikes per pixel and is robust to increasing occupancy. In experimental data, frames counted with BPC are shown to reconstruct diffraction peaks corresponding to individual nanoparticles with relatively higher intensity and produce images with improved contrast when compared to a standard counting method. Together, these results show that BPC excels in experiments where pixels see a high flux of electron irradiation such as in situ TEM movies and diffraction.

Improvements in electron microscopy (EM) techniques have been essential to advancing materials science^{1,2} and structural biology^{3–5} via imaging of objects at the nanoscale. Recent developments in direct electron detectors⁶ have increased the capabilities of electron microscopy, enabling improvements in imaging^{7,8}, diffraction^{9,10}, and spectroscopy¹¹. One type of detector, called an active pixel sensor (APS)¹², utilizes a complementary metal-oxide-semiconductor (CMOS) sensor with fast readout in order to reduce the number of electrons per frame to below about 1% of the pixel number. Thus, each electron strike produces a localized deposition of energy that can be differentiated from other electron strikes and essentially “counted” as one electron regardless of the energy deposited¹³. The process of converting an extended energy deposition to a single electron strike is known as electron counting. This process improves the signal-to-noise ratio relative to the raw sensor response and is central to EM experiments that involve direct electron detection based on the APS design^{14,15}. The “effective” detector quantum efficiency can be increased with improved electron counting algorithms, producing high signal-to-noise frames at a wider range of electron fluxes.

Counting electron strikes in APS detectors is challenging for several reasons. First, the amount of energy deposited in a thin volume of material by an energetic particle passing through it follows a Landau distribution¹⁶.

This captures both the most probable value of energy deposition, as well as a long tail extending to the full energy of the particle resulting from relatively rare interactions that leave behind large amounts of energy. For a large number of individual electron strikes in a given pixel the total energy deposition tends towards the most probable value, but for one to several hits, the Landau tail makes it difficult to quantify exactly the number of electrons that hit a pixel. Second, an electron can undergo multiple scattering in the medium, changing its direction of travel, producing clusters of energy responses with an inconsistent shape and extent¹⁷. Finally, there is electronic noise¹⁸ intrinsic to the CMOS hardware that can mask electron strikes with low energy deposition (false-negatives) or be incorrectly counted as electron strikes (false-positives). A counting algorithm that attempts to normalize the response to a single electron is therefore important in improving image quality. However, developing such an algorithm is challenging due to issues in segregating low-energy depositions from electronic noise and distinguishing multiple electron hits from single hits. These challenges are exacerbated by variations in the magnitude and shape of the energy deposition across the detector.

APS sensors are used to implement four-dimensional scanning transmission electron microscopy (4D-STEM)¹⁹, a mode of operation capable of mapping crystallographic information with atomic- to nanoscale

¹Instituto de Física Corpuscular (IFIC), CSIC & Universitat de València, Paterna, Valencia, Spain. ²Scientific Data Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA. ³Molecular Foundry, Lawrence Berkeley National Laboratory, Berkeley, CA, USA. ⁴Division of Molecular Biophysics & Integrated Bioimaging, Lawrence Berkeley National Laboratory, Berkeley, CA, USA. ⁵Engineering Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA.

✉e-mail: kebouchard@lbl.gov; percius@lbl.gov; agoldschmidt@lbl.gov

detail. In this technique, a focused electron probe is typically scanned over a sample in a two-dimensional (2D) raster pattern, though sparse and dynamic patterns are also possible, and a 2D diffraction pattern is measured at each probe position (see Fig. 1a). This results in a four-dimensional (4D) dataset with the potential to be hundreds of gigabytes to terabytes in size. 4D-STEM has become practical with the availability of fast, sensitive detectors leading to the need for computer hardware and algorithms to process such large datasets. To deal with this, electron counting algorithms are typically optimized for and built into specific camera hardware enabling in-line data processing during the experiment. Thus, a fast implementation of a counting algorithm is used (see ref. 20 for an open source implementation) which relies on simple thresholds to eliminate noise and energetic x-rays. These simple algorithms can only accommodate a low flux per frame on the camera and do not consider the shape of energy deposition over multiple pixels.

Other approaches to electron counting use artificial intelligence, which has been increasingly employed in electron microscopy in recent years^{21–24}. In particular, convolutional neural networks (CNN) have been studied to improve localization accuracy and reduce coincidence loss^{25,26}. However, the large numbers of parameters in these models, which must be optimized through back-propagation and stochastic gradient descent, require large amounts of simulated data and computation for training, and retraining is necessary for different detectors and electron energies. An approach that takes into account the energy deposition shape of individual electrons on the detector but does not require extensive detector/energy specific training could lead to substantial improvements that are broadly applicable, but are nascent¹³.

To enable more accurate nanoscale imaging, we developed a novel method called Back-Propagation Counting (BPC) for electron counting (see Fig. 1). The algorithm makes use of modern machine learning techniques and software²⁷ and back-propagation^{28,29} to fit the energy deposition of a raw 2D frame to a “counted” frame of the same dimensions based on the

expected average energy deposition of single electron strikes. The process does not involve training deep neural networks and thus does not require large amounts of training data. It can be readily applied with minimal understanding of the detector used in the acquisition. BPC processes low-occupancy frames with similar accuracy to standard methods and accommodates higher electron flux experiments with the ability to differentiate multiple electron hits per detector pixel. Through detailed characterization using synthetic and experimental data, we demonstrate that BPC results in quantitative and qualitative improvements in final data quality when compared to a standard counting method and removes artifacts. These results suggest that BPC offers improvements in materials imaging, especially in cases of high electron flux, enabling clearer reconstruction of material structures.

Results

Here we present electron counting results using the BPC method and compare them to a standard counting algorithm from *stempy*³⁰ (see Methods). We make use of both Monte Carlo simulation and experimental data from two different studies.

The BPC method counts electrons more consistently with increasing occupancy in synthetic data

A goal of BPC was to succeed at counting over a wide range of electron occupancies where the standard method fails. To validate that BPC achieved this goal, we constructed simulated frames using Monte-Carlo-generated electron hits. This allowed us to study BPC under controlled conditions of pixel response amplitude and occupancy. We consider a case in which the charge deposited by a single electron is far greater than the noise present in the frames so that the results are independent of effects due to noise, and we focus specifically on counting in regimes of different electron occupancy.

To study electron counting under conditions of varying occupancy, electron hit locations were placed according to a Gaussian distribution with




Fig. 1 | Illustration of 4D-STEM and Back-Propagation Counting (BPC). **a** A scanning electron nanobeam incident on a sample produces a diffraction pattern captured by a pixelated APS-based detector. Simultaneously, high-angle scatterers are captured by a monolithic high-angle annular dark field (HAADF) detector. **b** An estimated counted frame is computed from the raw frame and convolved with the

pre-computed 2D kernel to produce a reconstructed frame. Back-propagation is used to iteratively update the estimate according to a loss function. **c** The final optimized count grid is rounded to integer values. The ground-truth is shown for comparison but is not used in the algorithm.

a width of $\sigma = 3$ (note that this σ is the width of the Gaussian distribution of hit locations and is unrelated to σ_e from Eq. (1) of Methods), and the pixel responses for each hit were chosen randomly from the constructed database of simulated hit distributions (see Methods). A frame size of 51×51 pixels was chosen. A total of approximately 10,000 electron hits were distributed over N_{MC} frames with a number of electrons per frame selected from a Gaussian distribution with mean χ , such that $N_{MC} \cdot \chi = 10,000$, in (N_{MC}, χ) pairs of (10,000, 1), (1000, 10), (100, 100), and (10, 1000). The counting was performed per frame, so the final electron statistics were approximately the same, but the counting took place under conditions of different occupancy.

Figure 2 shows the simulated Monte Carlo Gaussian patterns and counting results for the different approximate numbers of electrons per frame. Figure 2a shows the sum of the true electron hits for all Monte-Carlo-generated frames. Note again that while different numbers of electrons were

incident on the frames shown in the four columns, the number of generated frames was varied so that the true number of hits was statistically similar in each case. Figure 2b, c show the counted hits summed over all frames for the BPC and standard counting methods respectively. BPC continues to count at higher pixel occupancy where the standard method breaks down. To more directly visualize this, Fig. 2d shows the projections of the various 2D plots on the y-axis, including the MC truth from row 2a, the standard method from row 2c, and several variations on the BPC method: assuming all pixels with nonzero electron count have exactly 1 ("unweighted"), the full electron counting with no attempt to correct for Landau fluctuations with the prior ("w/o prior"), and the full electron counting with the prior-based correction ("w/prior") from row 2b.

From these results, we conclude that the new BPC method outperforms the standard method at higher occupancy in synthetic data, even without




Fig. 2 | The BPC method counts electrons more consistently with increasing occupancy in synthetic data. A Gaussian ($\sigma = 3$ pixels) distribution of electron hits is simulated and counted for 1, 10, 100, and 1000 electrons per frame. The true number of electron hits summed over all frames in the Monte Carlo is shown in (a) for the different mean numbers of electron frames. The sums over all counted frames are shown using the BPC method in (b) and the standard method in (c). The final row

d shows the counts projected along the y-axis, summed over all frames, for: the Monte Carlo truth, the standard method, a version of the BPC method with a maximum of 1 electron per hit (unweighted), the BPC method without application of the Landau correction described in Methods (w/o prior), and the full BPC method (w/prior).




Fig. 3 | The BPC method consistently counts experimental data under a constant electron flux, evaluated over varying aperture sizes. The constant flux of electrons was measured over four different camera lengths to provide a stable number of electrons with different occupancies. Similar datasets at electron currents $I_0 \approx 30$ pA, $2I_0$, $4I_0$, and $8I_0$ were acquired. **a** The raw circular aperture datasets acquired, summed over 16,512 frames of 576×576 pixels, for electron current $8I_0$. **b** The total

number of electron counts summed over 16,512 frames vs. camera length for each of the 4 electron currents. The curves shown in all cases were normalized to the maximum number of counts obtained using the standard counting method. **c** Average difference in the number of counts per frame relative to the standard method for variations of the BPC method, and **d** variation [(max-min)/max] in the curves shown in (b) for each method, averaged over all electron currents.

employing weights to account for multiple electrons per hit. The exact numbers of electrons counted over all frames in each occupancy group, along with mean and RMS errors on the electrons counted/frame, are shown in the Supplementary Information (Table S1). As the “BPC w/o prior” variant allows for more than one counted electron per pixel but does not use the prior to account for Landau fluctuations, it consistently reports more counts. Also, the application of the prior-based correction has a greater effect at lower pixel occupancy, because the long tail of the effect of the Landau distribution is averaged out at higher occupancies.

The BPC method consistently counts experimental data under a constant electron flux, evaluated over varying aperture sizes

We next studied the performance of the BPC method under conditions of varying electron occupancy using experimental data. Unlike in Monte Carlo, we do not have the information of each true electron hit location. However, by using a fixed electron current and varying the post-specimen detector camera length, we can change electron density on the camera to test coincidence loss. Given the results from the simulation studies presented, we hypothesized that BPC would correctly count electrons at higher density.

Experimental data at different camera lengths was collected at 300 kV and a set of fixed beam currents. 16,512 frames of 576×576 pixels each were

acquired for four different microscope camera lengths: 47.3 mm, 59.9 mm, 74.3 mm, and 92 mm. This was repeated for each of four different electron currents: $I_0 \approx 30$ pA, $2I_0$, $4I_0$, and $8I_0$, for 16 datasets in total. Electron counting was performed using the standard method²⁰ and using the BPC method. The first 1000 frames were used in constructing the prior for the Landau fluctuation corrections.

The results of the counting experiment are summarized in Fig. 3. The top panel of the Fig. 3a shows the raw data acquired for the highest electron current (approximately 240 pA) and the four different camera lengths. The decreasing number of electrons per pixel with increasing camera length is evidenced by the decreasing intensity (color scale) from left-to-right. In Fig. 3b, for each electron current, the normalized sum of all counted electrons is shown for BPC and the standard method for the four different camera lengths (counts were normalized to the maximum number of counts over all camera lengths obtained with the standard method for each current). As in Fig. 2, three variations of the results of the BPC method are shown to better understand the effects of various components of the method. All curves corresponding to the BPC method showed substantially reduced variations in the total number of counted electrons vs. camera length when compared to the standard method. For the lower currents, the number of electrons reported between the BPC and standard methods is similar, and at higher

currents, BPC consistently reports higher count values. While the standard method does not account for an extended electron energy distribution or electron occupancy and assigns a weight of 1 to all counted pixels, BPC counts more consistently, as shown by the flatter normalized electron count vs. camera length curves even in the case in which one electron per counted pixel was assumed. This reflects an improved ability to count electrons that strike different nearby pixels, producing overlapping energy distributions.

These observations are further quantified in Fig. 3c, which shows the difference in the counts/frame between BPC and standard methods for each electron current studied, averaged over the different camera lengths. We found that BPC counts more electrons than the standard method on average as the current, and therefore electron occupancy, increases. Finally, Fig. 3d shows the average percentage variation of the counted electrons per frame. This is determined for each counting method by computing, for each electron current, the quantity $[(\max - \min) / \max]$, where $\max$ and $\min$ are the maximum and minimum counts per frame across all four camera lengths, and averaging this quantity over the 4 currents. As hypothesized, the percentage variation in the number of electrons counted at each camera length is reduced when employing the BPC algorithm, and in particular with the prior-based correction. This demonstrates the ability of BPC to more correctly count electrons crowded into a smaller region on the detector.

The BPC method enhances diffraction from nanoparticles, yielding higher intensity diffraction peaks and improved image clarity

Finally, we evaluated the performance of the method in identifying diffraction peaks corresponding to individual nanoparticles. Accurate reconstruction of single diffraction peaks, and in particular their intensities, is of critical importance in applications such as crystallography³⁰. Strong diffraction peaks are very likely to contain multiple electrons per frame, and thus the BPC method's ability to more accurately estimate the intensities of these peaks could lead to better crystallography metrics and structure solutions.

The dataset consisted of 262,144 frames of doped NaYF₄ nanoparticles (see Methods). A simultaneously acquired HAADF-STEM image was used to identify nanoparticles in real space. The raw frames, captured on the same detector as used in the studies with varying aperture, were counted with the standard method and the BPC method, employing 1000 frames to compute the occupancy prior for reduction of Landau over-counting. Figure 4 shows the summed diffraction patterns (Fig. 4a–c) and real space images showing the summed diffraction pattern intensity (excluding a region of radius 50 pixels about the direct beam) at each electron probe position (Fig. 4d–f), for which the counting was performed using the standard method (Fig. 4a, d) and the BPC method (Fig. 4b, e: unweighted, Fig. 4c, f: with occupancy weights and the application of the prior to correct for Landau fluctuations). We found that BPC yields visibly improved contrast and sharper diffraction peaks, even without employing occupancy weights, compared to the standard model.

The HAADF-STEM image acquired during the scan was used to identify individual nanoparticles in real space (see Methods). A diffraction pattern corresponding to a single nanoparticle is shown in Fig. 4g for the standard counting method and in Fig. 4h for the BPC counting method. Various electron diffraction peaks are visible in the patterns. A single diffraction peak is identified with an arrow in Fig. 4g, h and shown enlarged below the full diffraction patterns. A major effect that we found was that BPC eliminates a “halo”-like reconstruction artifact present in the standard method in regions of high intensity. This effect is due to the final step of the standard method, which only permits assigning a count to pixels with intensities greater than their nearest neighbors, thus disallowing adjacent counts and possibly creating halo-like voids in regions of high intensity.

We directly compared the intensities of single diffraction peaks by taking the difference $I_{\text{BPC}} - I_{\text{standard}}$ of the two patterns (Fig. 4i). We found that the diffraction peaks were in many cases equally or more intense using BPC. To quantify this observation, we extracted peaks from several patterns (see Methods), and 7725 peaks were identified across 970 nanoparticle diffraction patterns. Figure 4j shows the histogram of differences in intensities $I_{\text{BPC}} - I_{\text{standard}}$ for the common diffraction peaks that were found in both datasets.

Critically, we found that the vast majority of peak intensities were larger for BPC compared to the standard method, sometimes by large amounts.

Discussion

Direct electron detectors are commonly deployed to image materials at the nanoscale. We addressed the problem of electron counting for APS-based direct electron detectors by developing the BPC method. We showed that BPC counts electrons more consistently with increasing occupancy in synthetic data. Analogously, in experimental data, we demonstrated the ability of BPC to more correctly count the same number of electrons even when they are crowded into a smaller region on the detector. Finally, in nanobeam 4D-STEM diffraction experiments, we showed that the BPC method can produce diffraction patterns with visibly improved contrast and higher intensity diffraction peaks than the standard method. This is significant because the standard method is often used in 4D-STEM scans with APS detectors and was designed to work at low electron occupancy. Improved electron counting leads to more precise reconstructions of material properties—such as strain, composition, and local electromagnetic fields—critical for understanding nanoscale phenomena³¹. Thus, our work paves the way for improved characterization and understanding of materials.

While a thorough evaluation of quantities such as the modulation transfer function (MTF) and detector quantum efficiency (DQE) is beyond the scope of this study, we note that the application of a Gaussian convolution kernel allows BPC to utilize the full spatial information of an electron impact on a pixelated detector. As observed by McMullan et al.³² and others^{33–35}, a deterministic blur leads to an increase of DQE at higher spatial frequencies due to the suppression of aliasing artifacts associated with simple thresholding. While the BPC method still assigns each electron hit to a single pixel, it does account for the detector's PSF, which likely contributes to the observed improvement in image quality over the standard counting method.

The BPC method could potentially be improved in future work. While the BPC method produces good results for the 4D Camera, which is back-thinned³⁶, it may be less suitable for electron hits measured in thicker CMOS detectors, in which the deposited charge distribution deviates significantly from a Gaussian distribution. However, for certain thicker detectors such as hybrid pixel detectors (HPDs), the method may be applicable in cases where diffusion dictates the spreading of charge. Future extensions of the method could generalize the single-electron response kernel to a non-Gaussian form.

In conclusion, we developed a novel electron counting algorithm for APS detectors based on convolutions with a pre-parameterized Gaussian kernel and back-propagation. The algorithm showed improved performance in counting electrons, in particular under conditions of increased pixel occupancy. It requires only the selection of two parameters (A_e and σ_e), which can be estimated from a small detector calibration dataset. No data-intensive training step is necessary. While a faster standard counting method is most suitable for obtaining an initial first look at 4D-STEM data, the BPC method offers improvements, particularly in cases of higher electron occupancy as in the examples shown in this study, with significant implications for enhanced materials imaging.

Methods

The BPC method consists of the convolution of a Gaussian kernel over an estimated count grid, followed by optimization of the count grid values via back-propagation according to a squared error loss. It is essentially a deconvolution method implemented using back-propagation assuming a fixed response defined by the average single-electron energy deposition analogous to a point-spread function (PSF). A dataset of sparse electron hits on a given detector is helpful to characterize the average single-electron response, but the BPC method does not require a specific training phase or training data. Thus, the only important parameters that need to be specified can be estimated from either known detector characteristics (i.e. thickness, pixel size, etc.) or a relatively small amount of sparse electron data.

The standard counting method to which we compare BPC involves, in brief, the application of a lower threshold to eliminate background noise and




Fig. 4 | The BPC method enhances diffraction from nanoparticles, yielding higher intensity diffraction peaks and improved image clarity. The sum of counted diffraction patterns (a–c) and the real-space images constructed by summing the total diffraction pattern intensities at each probe position (d–f), as measured with a 4D-STEM detector and counted using the standard counting method and the BPC counting method, are shown. The diffraction patterns are shown with log-scale intensity (min. value = 4, max. value = 2932505), the real space images are shown with linear-scale intensity (min. value = 0, max. value = 618), and all images in each row (a–c) and (d–f) are shown with equal contrast. An example diffraction

pattern from the same dataset summed over real-space probe positions of a single nanoparticle, reconstructed with the standard (g) and BPC (h) counting methods, is also shown, as well as the difference $I_{\text{BPC}} - I_{\text{standard}}$ (i) between these two patterns. The images in (g, h) are shown on a logarithmic intensity scale (min. value = 0 assigned to the darkest intensity, max. value = 100). A single peak is highlighted to demonstrate a “halo” artifact observed in the standard counting for intense peaks. j Intensity difference for diffraction peaks selected from many nanoparticles, counted with the BPC and standard counting methods.

an upper threshold to eliminate x-rays. It is, to first order, similar to counting methods used in commercially available detectors. The method is described in more detail below (see “Standard electron counting with stempy”).

Figure 1 summarizes the BPC method workflow. Each electron is assumed to contribute its deposited charge in the shape of a 2D Gaussian with a specified amplitude and width. The parameters of the Gaussian

are dependent on the detector and primary electron energy and remain fixed. Initially, the raw frame (F) is converted to an estimated initial count grid (C^0 , see Fig. 1b), and the grid is optimized with stochastic gradient descent via back-propagation starting from this initial condition (described in more detail later). Each pixel in the count grid contains a floating-point value which is adjusted with each iteration in an attempt

to match the raw frame. Once the optimization is complete, the pixel counts are rounded to the nearest integer. Additionally, we can account for the Landau fluctuations of the electron energy deposition in cases where a single electron strike may be erroneously identified as several due to the Landau tail. The method was implemented in PyTorch²⁷. AI models including ChatGPT and Google Gemini were employed to aid in the construction of the code and corresponding analysis, and all results from the use of AI were verified by the authors. Each step of the algorithm is described in more detail below.

Single-electron Gaussian parametrization

The BPC method first requires the selection of appropriate parameters for the 2D Gaussian kernel, which estimates the single-electron energy deposition signature at the incident energy of interest. The kernel size $W \times W$ (W odd) must first be chosen. In this study, we use $W = 3$ though the implementation of the method makes it straightforward to increase to $W = 5$ or greater. An approximate determination of the amplitude A_e and width σ_e of the kernel can be extracted from a low electron flux dataset in which the vast majority of electron strikes are single-occupancy. Starting from the raw data in ADC counts and subtracting the baseline, a nominal single-electron threshold can be set, and patches of size $W \times W$, centered on pixels exceeding this threshold, can be averaged and fit to a 2D Gaussian,

$$G(A_e, \sigma_e) = A_e \exp\left(-\frac{x^2 + y^2}{2\sigma_e^2}\right). \quad (1)$$

The choice of parameters is not bound by such a fit, and in practice the parameters obtained may not be optimal. For the experimental data analyzed in this study, we use $A_e = 19.03$ and $\sigma_e = 0.51$ which were extracted from a 2D Gaussian fit to an average single-electron hit distribution in the sparsest dataset (with lowest current) analyzed in the studies of varying aperture size. In the synthetic data analyzed, a similar 2D Gaussian fit to frames constructed with a sparse electron distribution gave $A_e = 28$ and $\sigma_e = 0.77$. We reduced σ_e to 0.6, as the algorithm was found to be unstable (converged slowly and produced spurious counts) for values of σ_e significantly greater than 0.6.

Electron counting fit with back-propagation

Next, BPC reconstructs each raw frame by applying the 2D Gaussian kernel to a “count grid” where each pixel estimates a number of electron strikes. An initial count grid C_{ij}^0 is first estimated from the raw frame F by subtracting a baseline b based on the average dark noise and normalized by A_e .

$$C_{ij}^0 = \left\lfloor \frac{F_{ij} - b}{A_e} \right\rfloor. \quad (2)$$

Here, $\lfloor \cdot \rfloor$ refers to rounding to the nearest integer. The 2D Gaussian kernel is then convolved with the count grid to obtain a reconstructed frame,

$$\begin{aligned} P_{ij}^k &= C_{ij}^k * G(A_e, \sigma_e) \\ &= \sum_{l,m} C_{(i+l)(j+m)}^k G_{lm}(A_e, \sigma_e), \end{aligned} \quad (3)$$

where l and m range from $-(W-1)/2$ to $(W-1)/2$. The elements of the count grid C_{ij}^k are trainable parameters that serve as “weights” for the Gaussian kernel, and we denote the predicted frame as P_{ij}^k and the measured frame as F_{ij} . We then compute the loss for the iteration, which we choose to be the sum of the squared differences of the pixels in the true and reconstructed frames,

$$L_k = \sum_{i,j} (F_{ij} - P_{ij}^k)^2. \quad (4)$$

We then perform back-propagation with learning rate λ (in this work we use PyTorch’s implementation of the Adam optimizer with base $\lambda =$

0.01) to converge on a count grid that best reproduces the raw frame (Fig. 1b),

$$C^{k+1} = C^k - \lambda \frac{\partial L_k}{\partial C^k}. \quad (5)$$

By back-propagating the gradients through the loss L_k , we update the count grid for the next iteration to obtain C_{ij}^{k+1} . This strategy of updating the “input” via back-propagation rather than the parameters of the convolutional kernel is similar to that employed in ref. 37 to produce images matching specific styles. Input optimization for a fixed model can also be found in studies of adversarial examples³⁸ in neural networks.

The procedure can be stopped after a given number of iterations M . Increasing M results in a counted grid that more precisely matches the experimental data but obtaining a perfect match is not required for good counting. For the nanoparticle counting experiment, we set a maximum of 3000 iterations. A dynamic stopping condition can also be implemented to activate when the loss decreases by less than a specified percentage over a set number of consecutive iterations. The final count grid is rounded to the nearest integer to obtain an integer number of electron counts for each pixel.

Addressing the Landau tail with a prior

As described earlier, the energy deposited by an electron in a thin detector will follow a Landau distribution. The long tail of this distribution indicates that a single electron can deposit, with relatively high frequency, significantly more energy than the most probable deposited energy. This will likely lead to many pixels being assigned high electron counting numbers, as BPC assumes an energy deposition shape with a fixed amplitude that will most likely be chosen to be within a factor of ~ 2 of the most probable value in the Landau curve.

It is not possible to directly determine whether a given pattern of energy deposition belongs to one electron depositing much more energy than the most probable value or multiple electrons each depositing much less energy; however, we can use a statistical method to modify the optimized count grid in a way that aims to improve the counting. The procedure is as follows:

1. Construct a “prior” μ_{ij} which estimates the average number of electrons per frame striking each pixel in the counted dataset.
2. Use the prior to compute the probability that, if a pixel is hit, it was hit by more than one electron.
3. For all pixels in the optimized count grid C_{ij} containing more than 1 hit count, reduce the hit count to 1 if a uniformly distributed random number is greater than the probability computed in step #2.

Note that the procedure makes no assumptions about the actual Landau distribution of energy deposition but rather removes cases of multiple hits when, based on the expected number of hits and Poisson statistics, they would be considered rare. Here, we assume that the electron occupancy across frames remains constant throughout the scan. As this is rarely true in practice but is often a good approximation, in this study we show counting results with and without the application of this procedure for comparison.

The prior μ_{ij} can be estimated from a number N_p of frames F_m as the average of these frames divided by a chosen characteristic single-electron amplitude. The value A_e from equation (1) can be used as this amplitude, and we require a minimum mean value of zero. Therefore, we have,

$$\mu_{ij} = \max\left\{0, \left(\frac{1}{N_p} \sum_{m=1}^{N_p} (F_m)_{i,j}\right) / A_e\right\}. \quad (6)$$

Once the prior has been constructed, we take it to represent the mean number of electrons per frame striking each pixel. Assuming that this number follows a Poisson distribution, we can use it to compute the probability that n_{ij} electrons strike pixel (i, j) as $\text{Poisson}(n_{ij}; \mu_{ij})$. The quantity we are interested in for the purposes of reducing the counting error due to

the Landau tail is the conditional probability that, given at least 1 electron hit, we have two or more hits, $P(n_{ij} \geq 2 | n_{ij} \geq 1; \mu_{ij})$. According to Bayes' theorem,

$$\begin{aligned} P(n_{ij} \geq 2 | n_{ij} \geq 1; \mu_{ij}) &= \frac{P(n_{ij} \geq 1 | n_{ij} \geq 2; \mu_{ij}) \cdot P(n_{ij} \geq 2; \mu_{ij})}{P(n_{ij} \geq 1; \mu_{ij})} \\ &= \frac{P(n_{ij} \geq 2; \mu_{ij})}{P(n_{ij} \geq 1; \mu_{ij})} \end{aligned} \quad (7)$$

since $P(n_{ij} \geq 1 | n_{ij} \geq 2; \mu_{ij}) = 1$. We can compute this conditional probability, knowing that the probability of drawing a Poisson number greater than some integer l is equal to 1 minus the sum of all Poisson probabilities for numbers less than l , that is,

$$\begin{aligned} P(n_{ij} \geq l; \mu_{ij}) &= 1 - \sum_{m=0}^{l-1} \text{Poisson}(m; \mu_{ij}) \\ &= 1 - e^{-\mu_{ij}} \sum_{m=0}^{l-1} \frac{(\mu_{ij})^m}{m!}. \end{aligned} \quad (8)$$

With this, we compute $P(n_{ij} \geq 2 | n_{ij} \geq 1; \mu_{ij})$ and compare it to a random number generated from a uniform distribution in the $[0,1]$ interval for each pixel. For those pixels with more than 1 electron hit in the final counted grid C_{ij} , if the generated random number is greater than the computed probability, we set the value for that pixel to 1. In this way, improbable multiple hits are reduced to a single hit which was assumed to have deposited more energy due to the Landau tail. But for pixels with a high average number of electron hits, multiple hits are more likely to remain in the final counted grid.

This procedure does not address all possibilities. It simplifies the decision for each pixel to “single-electron” or “multi-electron”. It also makes no attempt to determine whether a pixel with 3 apparent electron hits is actually 2 hits with 1 hit depositing above-average energy. A more sophisticated procedure would need to take into account the possible combinatorics of the composition of an n -electron hit response and the details of the Landau distribution, a task which is NP-hard and infeasible to confront in a practical application.

Practical details of the BPC method

For the nanoparticle dataset (Fig. 4), we processed batches of 2064 frames each of which required approximately 5 minutes to reach 3000 iterations using an NVIDIA A100 GPU. The total processing time was 10.6 hours. The method is slower than threshold-based counting methods, and is better used for post-processing. The computational cost scales linearly with the number of pixels and iterations. While a single iteration is independent of electron occupancy, the occupancy and content of a frame affect the number of iterations required for convergence. Thus, more complex frames require more iterations to converge.

Standard electron counting with stempy

The standard electron counting algorithm used in this study is implemented in the stempy²⁰ 4D-STEM analysis package. It takes as inputs the raw detector frame, a background threshold, and an X-ray threshold to exclude high-intensity events that were likely to have been the result of X-rays striking the detector. A dark reference frame for background subtraction and a gain map for pixel-wise sensitivity correction can also be provided as optional inputs.

The algorithm first applies the dark reference subtraction and gain correction (if provided) to each frame. The pixel intensities are then thresholded, so that pixels with values below the specified background threshold or above the X-ray threshold are excluded from consideration as electron hits. An electron strike is registered at a pixel's location if its intensity is greater than that of all of its eight immediate neighbors. In our analysis of simulated electron strikes, a background threshold of 10 counts and an X-ray threshold of 175 counts were employed.

Simulation of electron hits in APS detectors

To construct a simulated frame containing some number of electron hits, we first generate a library of pixel response distributions for single-electrons incident on a silicon plate using Geant4³⁹, which models the Landau nature of the electron energy deposition. For each electron hit in the simulated frame, we can then place a randomly-selected pixel response distribution centered on its specific hit location. In this study, we used an incident electron energy of 300 keV and an active region thickness of 5 μm , and the electron hit locations on a single frame were simulated according to a Gaussian distribution.

To model the charge collection process, diffusion was applied from the energy deposition sites to the edge of the plate. This was done by simulating a random walk for each electron produced in the Geant4 simulation using an estimated mean electron scattering length of 1 μm . The detection plane was binned to a $10 \times 10 \mu\text{m}^2$ pixel size, and the simulated pixel counts were then multiplied by a scaling factor so that the peak of the Landau distribution of the sum of all ADC counts in a single hit was approximately 40. 200,000 single-electron hits were saved in a database for use in building multi-hit frames. In the simulation studies presented here, we consider frames for which electronic noise is insignificant, and therefore we add a low Gaussian noise of mean 0 and width of 1 to all pixels in the simulated frames.

Experimental details

Near parallel-beam 4D-STEM datasets were acquired using the TEAM 0.5 aberration corrected STEM⁴⁰ of the National Center for Electron Microscopy facility of the Molecular Foundry at Lawrence Berkeley National Laboratory. The STEM was operated at an accelerating voltage of 300 kV for all experiments. An APS detector called the 4D Camera⁴¹ was used for the data collection with a frame rate of 87,000 Hz and 576 by 576 pixels per frame. The data shown in Fig. 3a were generated using a constant beam current (flux) of approximately 30 pA, and the post-specimen camera length was reduced to increase the current density on the camera. The process was also repeated with higher flux values to increase the likelihood of coincidence on the camera. The scan of nanoparticles presented in Fig. 4 was acquired with a probe convergence angle of approximately 0.1 mrad and about 15 pA of beam current. 512 by 512 real space probe positions were used to image the 614 nm field of view with a probe step size of 1.2 nm. The size of the resulting raw dataset was 174 GB which was processed using the standard counting method and our newly developed BPC method.

The nanoparticle samples consisted of hexagonal-phase NaYF_4 core/shell nanoparticles with 15% Tm^{3+} core and 20% Gd^{3+} shell doping^{42,43}. These Tm^{3+} -based nanoparticles can both upconvert short-wave infrared light to shorter near-infrared wavelengths⁴³ for use in imaging biological systems and be used as remote force sensors with exceptionally large dynamic range⁴⁴. Their single-crystalline nature produces very strong Bragg peaks in diffraction patterns with a high probability of multiple electron strikes in detector frames.

For diffraction peak finding and fitting we used functions from the ncempy Python library⁴⁵. The center of each nanoparticle was identified in the simultaneously acquired HAADF-STEM image (see Supplementary Information, Fig. S1), and a 2D circular mask of radius 6 pixels was used to extract each particle's real-space coordinates. We sum all diffraction patterns corresponding to the masked area to generate one diffraction pattern for each individual nanoparticle.

Bragg peaks were found in each summed diffraction pattern using a common peak finding algorithm employing a threshold value of 12 electrons to remove background noise. Only the most intense peak was kept when multiple peaks were identified within a minimum distance of 10 pixels, and only peaks common to both the BPC and standard reconstructions were considered. The peaks were fit to 2D Gaussians using an 11×11 pixel window centered on the peak location as the region of interest. The intensity of each peak was taken to be the maximum value obtained from the 2D Gaussian fit.

Data availability

The 4D-STEM datasets analyzed during the current study are not publicly available due to their large size but are available from the corresponding authors on reasonable request.

Code availability

The code for the BPC method is available on GitHub at: <https://github.com/jerenner/backpropcount>.

Received: 19 July 2025; Accepted: 29 April 2026;

Published online: 12 May 2026

References

- Ophus, C. Quantitative scanning transmission electron microscopy for materials science: Imaging, diffraction, spectroscopy, and tomography. *Annu. Rev. Mater. Res.* **53**, 105–141 (2023).
- Lin, Y. et al. Analytical transmission electron microscopy for emerging advanced materials. *Matter* **4**, 2309–2339 (2021).
- Gallagher-Jones, M. et al. Nanoscale mosaicity revealed in peptide microcrystals by scanning electron nanodiffraction. *Commun. Biol.* **2**, 26 (2019).
- Lazić, I. et al. Single-particle cryo-EM structures from iDPC-STEM at near-atomic resolution. *Nat. Methods* **19**, 1126–1136 (2022).
- Zhou, L. et al. Low-dose phase retrieval of biological specimens using cryo-electron ptychography. *Nat. Commun.* **11**, 2773 (2020).
- MacLaren, I., Macgregor, T. A., Allen, C. S. & Kirkland, A. I. Detectors—The ongoing revolution in scanning transmission electron microscopy and why this important to material characterization. *APL Mater.* **8**, 110901 (2020).
- Chen, Z. et al. Electron ptychography achieves atomic-resolution limits set by lattice vibrations. *Science* **372**, 826–831 (2021).
- Nakane, T. et al. Single-particle cryo-EM at atomic resolution. *Nature* **587**, 152–156 (2020).
- Hattne, J., Clabbers, M. T., Martynowycz, M. W. & Gonen, T. Electron counting with direct electron detectors in MicroED. *Structure* **31**, 1504–1509.e1 (2023).
- Saha, A. et al. Imaging the progression of radiolytic damage in molecular crystals with scanning nanobeam electron diffraction. *Phys. Rev. Lett.* **134**, 146101 (2025).
- Hart, J. L. et al. Direct detection electron energy-loss spectroscopy: a method to push the limits of resolution and sensitivity. *Sci. Rep.* **7**, 8243 (2017).
- Denes, P., Bussat, J.-M., Lee, Z. & Radmilovic, V. Active pixel sensors for electron microscopy. *Nucl. Instrum. Meth. A* **579**, 891–894 (2007).
- Battaglia, M., Contarato, D., Denes, P. & Giubileo, P. Cluster imaging with a direct detection CMOS pixel sensor in transmission electron microscopy. *Nucl. Instrum. Meth. A* **608**, 363–365 (2009).
- Zhu, D., Shi, H., Wu, C. & Zhang, X. An electron counting algorithm improves imaging of proteins with low-acceleration-voltage cryo-electron microscope. *Commun. Biol.* **5**, 321 (2022).
- Peters, J. J. P. et al. Electron counting detectors in scanning transmission electron microscopy via hardware signal processing. *Nat. Commun.* **14**, 5184 (2023).
- Landau, L. D. On the energy loss of fast particles by ionization. *J. Phys.* **8**, 201 (1944).
- Navas, S. et al. Review of particle physics. *Phys. Rev. D* **110**, 030001 (2024).
- McMullan, G., Faruqi, A., Clare, D. & Henderson, R. Comparison of optimal performance at 300keV of three direct electron detectors for use in low dose electron microscopy. *Ultramicroscopy* **147**, 156–163 (2014).
- Ophus, C. Four-dimensional scanning transmission electron microscopy (4D-STEM): from scanning nanodiffraction to ptychography and beyond. *Microsc. Microanal.* **25**, 563–582 (2019).
- Kitware. stempy: Toolkit for 4D STEM processing on HPC using a Python interface and C++ engine. <https://github.com/OpenChemistry/stempy>.
- Botifoll, M., Pinto-Huguet, I. & Arbiol, J. Machine learning in electron microscopy for advanced nanocharacterization: current developments, available tools and future outlook. *Nanoscale Horiz.* **7**, 1427–1477 (2022).
- Treder, K. P., Huang, C., Kim, J. S. & Kirkland, A. I. Applications of deep learning in electron microscopy. *Microscopy* **71**, i100–i115 (2022).
- Gleason, S. P., Lu, D. & Ciston, J. Prediction of the Cu oxidation state from EELS and XAS spectra using supervised machine learning. *npj Comput. Mater.* **10**, 221 (2024).
- Lobato, I., Friedrich, T. & Van Aert, S. Deep convolutional neural networks to restore single-shot electron microscopy images. *npj Comput. Mater.* **10**, 10 (2024).
- van Schayck, J. P. et al. Sub-pixel electron detection using a convolutional neural network. *Ultramicroscopy* **218**, 113091 (2020).
- Wei, J. et al. Deep learning approach for high-accuracy electron counting of monolithic active pixel sensor-type direct electron detectors at increased electron dose. *Microsc. Microanal.* **29**, 2026–2036 (2023).
- Paszke, A. et al. PyTorch: an imperative style, high-performance deep learning library. *Adv. Neural Inform. Process. Syst.* **32**, 8026–8037 (2019).
- Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
- Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning*. <http://www.deeplearningbook.org>. (MIT Press, 2016).
- Borek, D., Minor, W. & Otwinowski, Z. Measurement errors and their consequences in protein crystallography. *Acta Crystallogr. Section D* **59**, 2031–2038 (2003).
- Chen, Z. et al. Lorentz electron ptychography for imaging magnetic textures beyond the diffraction limit. *Nat. Nanotechnol.* **17**, 1165–1170 (2022).
- McMullan, G., Chen, S., Henderson, R. & Faruqi, A. Detective quantum efficiency of electron area detectors in electron microscopy. *Ultramicroscopy* **109**, 1126–1143 (2009).
- Paton, K. A. et al. Quantifying the performance of a hybrid pixel detector with GaAs:cr sensor for transmission electron microscopy. *Ultramicroscopy* **227**, 113298 (2021).
- van Schayck, J. P., Zhang, Y., Knoops, K., Peters, P. J. & Ravelli, R. B. G. Integration of an event-driven timepix3 hybrid pixel detector into a cryo-em workflow. *Microsc. Microanal.* **29**, 352–363 (2022).
- Zambon, P. Enhanced DQE and sub-pixel resolution by single-event processing in counting hybrid pixel electron detectors: a simulation study. *Front. Phys.* **11**, 1123787 (2023).
- McMullan, G. et al. Experimental observation of the improvement in MTF from backthinning a CMOS direct electron detector. *Ultramicroscopy* **109**, 1144–1147 (2009).
- Gatys, L. A., Ecker, A. S. & Bethge, M. A neural algorithm of artistic style <https://arxiv.org/abs/1508.06576>. (2015).
- Goodfellow, I. J., Shlens, J. & Szegedy, C. Explaining and harnessing adversarial examples <https://arxiv.org/abs/1412.6572>. (2015).
- Agostinelli, S. et al. Geant4—a simulation toolkit. *NIMA* **506**, 250–303 (2003).
- Dahmen, U. et al. An update on the TEAM project — first results from the TEAM 0.5 microscope, and its future development. in *EMC 2008 14th European Microscopy Congress 1–5 September 2008, Aachen, Germany*, (eds Luysberg, M., Tillmann, K. & Weirich, T.) 3–4 (Springer Berlin Heidelberg, Berlin, 2008).
- Ercius, P. et al. The 4D camera: An 87 kHz direct electron detector for scanning/transmission electron microscopy. *Microsc. Microanal.* **30**, 903–912 (2024).

42. Pedroso, C. C. S. et al. Immunotargeting of nanocrystals by spycatcher conjugation of engineered antibodies. *ACS Nano* **15**, 18374–18384 (2021).
43. Qi, X. et al. Short-wave infrared upconverting nanoparticles. *J. Am. Chem. Soc.* **146**, 29292–29296 (2024).
44. Fardian-Melamed, N. et al. Infrared nanosensors of piconewton to micronewton forces. *Nature* **637**, 70–75 (2025).
45. openNCEM. <https://github.com/ercius/openNCEM>.

Acknowledgements

This work was primarily funded by the US Department of Energy in the program "4D Camera Distillery: From Massive Electron Microscopy Scattering Data to Useful Information with AI/ML." JR acknowledges support from the Generalitat Valenciana of Spain under grant CIDEXG/2023/16. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a Department of Energy Office of Science User Facility using NERSC award (BES-ERCAP0032504). Work at the Molecular Foundry was supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This work was in part supported by the Laboratory Directed Research and Development Program of Lawrence Berkeley National Laboratory under U.S. Department of Energy Contract No. DE-AC02-05CH11231.

Author contributions

J.R. and A.G. devised the BPC counting method and J.R. implemented it in Python. J.R. wrote the initial draft of the manuscript and led revisions. A.G., J.R., K.B., M.W., and P.E. identified test cases for the method and contributed to the manuscript. M.W. provided PyTorch and ML expertise. P.E. devised the experiments carried out in the study and A.S. acquired the experimental data. C.P. and B.C. provided the nanoparticle samples. P.D. supervised the research. All team members contributed to the final revision of the article.

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/s41524-026-02126-6>.

Correspondence and requests for materials should be addressed to Kristofer Bouchard, Peter Ercius or Azriel Goldschmidt.

Reprints and permissions information is available at <http://www.nature.com/reprints>

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-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, 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 you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. 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-nc-nd/4.0/>.

© The Author(s) 2026