

Solving structure with sparse, randomly-oriented x-ray data

Hugh T. Philipp,^{1,*} Kartik Ayer,¹ Mark W. Tate,¹ Veit Elser,¹ and Sol M. Gruner^{1,2}

¹Cornell University, Laboratory of Atomic and Solid State Physics, Ithaca, NY USA

²Cornell's High Energy Synchrotron Source (CHESS), Ithaca, NY USA

*htp2@cornell.edu

Abstract: Single-particle imaging experiments of biomolecules at x-ray free-electron lasers (XFELs) require processing hundreds of thousands of images that contain very few x-rays. Each low-fluence image of the diffraction pattern is produced by a single, randomly oriented particle, such as a protein. We demonstrate the feasibility of recovering structural information at these extremes using low-fluence images of a randomly oriented 2D x-ray mask. Successful reconstruction is obtained with images averaging only 2.5 photons per frame, where it seems doubtful there could be information about the state of rotation, let alone the image contrast. This is accomplished with an expectation maximization algorithm that processes the low-fluence data in aggregate, and without any prior knowledge of the object or its orientation. The versatility of the method promises, more generally, to redefine what measurement scenarios can provide useful signal.

© 2012 Optical Society of America

OCIS codes: (040.7480) X-rays, soft x-rays, extreme ultraviolet (EUV); (110.7440) X-ray imaging; (000.2190) Experimental physics; (110.3055) Information theoretical analysis; (110.4155) Multiframe image processing; (040.0040) Detectors.

References and links

1. R. Neutze, R. Wouts, D. van der Spoel, E. Weckert, and J. Hajdu, "Potential for biomolecular imaging with femtosecond x-ray pulses," *Nature* **406**, 752–757 (2000).
2. V. Elser, "Noise limits on reconstructing diffraction signals from random tomographs," *IEEE Trans. Inf. Theory* **55**, 4715–4722 (2009).
3. N.-T. D. Loh and V. Elser, "Reconstruction algorithm for single-particle diffraction imaging experiments," *Phys. Rev. E* **80**, 026705 (2009).
4. N. D. Loh, M. J. Bogan, V. Elser, A. Barty, S. Boutet, S. Bajt, J. Hajdu, T. Ekeberg, F. R. N. C. Maia, J. Schulz, M. M. Seibert, B. Iwan, N. Timneanu, S. Marchesini, I. Schlichting, R. L. Shoeman, L. Lomb, M. Frank, M. Liang, and H. N. Chapman, "Cryptotomography: reconstructing 3D Fourier intensities from randomly oriented single-shot diffraction patterns," *Phys. Rev. Lett.* **104**, 225501 (2010).
5. H. T. Philipp, L. J. Koerner, M. S. Hromalik, M. W. Tate, and S. M. Gruner, "Femtosecond radiation experiment detector for x-ray free-electron laser (XFEL) coherent x-ray imaging," *IEEE Trans. Nucl. Sci.* **57**, 3795–3799 (2010).
6. H. T. Philipp, M. W. Tate, and S. M. Gruner, "Low-flux measurements with Cornell's LCLS integrating pixel array detector," *J. Inst.* **6**, C11006 (2011).
7. L. E. Baum, T. Petrie, G. Soules, and N. Weiss, "A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains," *Ann. Math. Statist.* **41**, 164–171 (1970).
8. G. Huldt, A. Szoke, and J. Hajdu, "Diffraction imaging of single particles and biomolecules," *J. Struct. Biol.* **144**, 219–227 (2003).
9. R. R. Coifman, Y. Shkolnisky, F. J. Sigworth, and A. Singer, "Graph Laplacian tomography from unknown random projections," *IEEE Trans. Image Proc.* **17**, 1891–1899 (2008).

10. D. Giannakis, P. Schwander, and A. Ourmazd, "The symmetries of image formation by scattering. I. Theoretical framework," arXiv:1009.5035 (2010).
 11. P. Thibault, M. Dierolf, A. Menzel, O. Bunk, C. David, and F. Pfeiffer, "High-resolution scanning x-ray diffraction microscopy," *Science* **321**, 379–382 (2008).
 12. D. H. Bilderback, J. D. Brock, D. S. Dale, K. D. Finkelstein, M. A. Pfeifer, and S. M. Gruner, "Energy recovery linac (ERL) coherent hard x-ray sources," *New J. Phys.* **12**, 035011 (2010).
 13. C. E. Shannon, "A Mathematical Theory of Communication," *Bell Syst. Tech. J.*, **27**, 379–423, 623–656 (1948).
-

1. Introduction

Ultra-intense, ultra-fast x-ray pulses from x-ray free electron lasers (XFELs), such as the Linac Coherent Light Source (LCLS), hold potential to provide structural information about proteins for which crystals are unavailable [1]. This so-called single particle imaging experiment involves scattering x-rays off many copies of a given protein, one protein at a time. This yields a sequence of x-ray detector images, each resulting from the diffraction of a single pulse scattered off a single protein. Since only some thousands of x-rays are scattered off each protein, each x-ray image, which consists of a few million pixels, is very sparsely populated with data: On average, each pixel in each image receives far fewer than one x-ray. If each protein intersected the x-ray beam in the same orientation, one could simply add many images to recover a statistically significant data set. However, each protein intersects the x-ray pulses in unknown, random orientations and there are too few x-rays in any single image to determine the orientation of that protein. The challenge is to devise a method that combines information from a large number of these severely Poisson noise limited images into a complete, consistently oriented data set [2].

A data reduction method has been proposed [3, 4] that should work, in principle, even in the case where the number of x-rays per image is so small that it is impossible to discover similar orientations by cross-correlating images. However, this method has not been experimentally tested in the extreme case of only a few photons per image. Here, we demonstrate that a simplified 2D version of the algorithm is capable of reducing this kind of data even with images that average only 2.5 x-ray photons per frame ($\sim 10^{-4}$ photons per pixel per frame). For this demonstration, in order to emulate realistic detector noise performance, we use the same pixel array detector chip that makes up the detector installed at the LCLS Coherent X-ray Imaging (CXI) beamline for the protein imaging experiment [5, 6]. The CXI detector differs from the one used here in that the former is an array of these chips to cover a larger area.

Our experiment simulates the diffraction patterns of randomly rotated particles by means of a mask placed in front of the detector that is given random rotations and is uniformly illuminated by a highly attenuated beam. We argue that this experiment contains all the salient features of a full 3D intensity reconstruction, namely unknown sample orientations and very low signal. The additional features of the single particle imaging experiment are the extra degrees of rotational freedom of the sample in 3D and the increased size of the detector. Since we discretize the space of rotations, the additional degrees of freedom only increases the size of the discrete set of orientations required to adequately sample the space. This does not make a qualitative difference to the problem. Similarly, a larger detector only increases the computational resources required. The reconstruction algorithm scales linearly as the number of orientational samples and the total number of photons imaged. Simulations have shown that the higher computational requirements can be handled comfortably [3].

2. Expectation maximization algorithm

The algorithm for reconstructing the x-ray intensity from data follows the expectation maximization (EM) principle [7]. EM starts with a random model of the intensity $w(i)$, where each pixel i is assigned a random value uniformly in the range $[0, 1]$. These values are iteratively up-




Fig. 1. (a) The lead x-ray mask mounted within an aperture in an aluminum disk. (b) A static x-ray image of the pattern collected as 432 individual frames with approximately 1/5 photon per pixel per frame. The frames were thresholded and averaged. (c) A reconstruction using randomly-oriented data having an average 11.5 photons/frame and 1.2 million recorded photons. (d) A reconstruction using randomly-oriented data having an average 2.5 photons/frame and 1.2 million recorded photons ([Media 1](#)).

dated by a rule that can only increase the likelihood of the model. The initial model is random because no information about the model is known.

Each iteration involves two steps. In the first step, each frame of data, f , is assigned a probability distribution, $p_f(r)$, with respect to its unknown rotation, r , relative to the current intensity model. The rotations are sampled in increments of $2\pi/N$, where N defines the angular resolution of the reconstruction. Each frame comprises photon occupancy, $k_f(i)$, at pixel i , which in our low-fluence experiment are almost all zero, the exceptions being equal to 1. Because the photon counts are independent Poisson samples of the intensity at each pixel, the probability is

$$p_f(r) \propto \prod_i \frac{w(i+r)^{k_f(i)}}{k_f(i)!} e^{-w(i+r)} \propto \prod_{i \in I_f} w(i+r), \quad (1)$$

where $i+r$ is rotation r , applied to pixel i , I_f is the set of pixels recording photons in frame f , with the final expression applying in the low-fluence limit. After normalizing the distributions, $p_f(r)$, the algorithm proceeds to the second step.

In the second step the algorithm aggregates the photon data from all the frames according to the distributions obtained in the first step:

$$w'(i) = \left\langle \sum_r p_f(r) k_f(i-r) \right\rangle_f. \quad (2)$$

The updated intensity model $w'(i)$ is an average of the photon counts in all frames with the appropriate distribution of rotations applied to each one. Linear interpolation is used in both steps, when mapping one square grid (model) onto one that has been rotated (detector). We consider the reconstruction to have converged when the root-mean-square difference between successive models is below a cutoff, chosen to be 0.0001 times the mean pixel value. The EM algorithm is still valid when the data is winnowed by a structure-neutral criterion, such as the photon occupancy. In our implementation, for example, we discarded all frames with zero occupancy.

An alternative approach being considered [8] for determining the rotations of randomly oriented particles involves classifying data on the basis of cross-correlations:

$$c_{ff'} = \sum_i k_f(i)k_{f'}(i). \quad (3)$$

This is not viable in the ultra-low fluence limit because the numbers $c_{ff'}$ are essentially all zero, and in any event do not distinguish frames derived from like or unlike particle orientations. A proposal [9, 10] to orient data by means of diffusion modes on a graph is inapplicable in these circumstances: in our experiment the associated graph comprises nearly a million nodes (one per frame) but only a handful of edges (rare instances where $c_{ff'} > 0$). The EM algorithm, by contrast, compares each frame not with other frames but with a model. A greater sensitivity of rotation determination in the EM algorithm can be traced to the multiplicative nature of the comparison expressed by Eq. (1).

The EM algorithm should in principle work with data of arbitrarily low-fluence. It is clear that this is the case when we consider that there will be rare fluctuations where the photon occupancy is 2 or greater, even when the mean is just a fraction of a photon. A fair assessment of the viability of reconstructions in the low-fluence regime must therefore take into account the inevitability of background. The effects of background in the interpretation of our results are well captured by a simple information rate ratio:

$$R = \frac{\sigma(1 + \text{SN}^{-1}) \log(1 + \text{SN}) - (\sigma + \text{SN}^{-1}) \log(1 + \sigma \text{SN})}{-\sigma \log \sigma}. \quad (4)$$

This is the information rate at signal-to-noise SN (the ratio of signal to background fluence) divided by the rate in the absence of background, where σ is the fraction of uncovered pixels.

Equation (4) is based on prior distributions for both the signal and background; a derivation is given in the appendix. The signal for our experiment was modeled as two-valued, corresponding to zero for pixels covered by the mask and a constant nonzero value for uncovered pixels. A single parameter, the fraction σ of uncovered pixels, describes this signal prior. When modeling a true diffraction signal, the prior would instead be the Wilson distribution. To model the background, we used a single-valued distribution corresponding to uniform background counts across the detector.

As an example of the application of Eq. (4), consider the case of $\text{SN} = 1/10$, which for $\sigma = 0.6$ gives $R \approx 0.01$. A low fluence experiment with 2 signal photons per frame and this poor signal-to-noise would therefore be like a zero-background experiment with only $2R = 0.02$ photons per frame. Applying Poisson statistics to this low rate we find that only about 1 in 5000 frames would have 2 or more photons and be actually useful for the reconstruction.

3. Data collection

The mask for our experiment was cut out of an x-ray opaque lead sheet (Fig. 1(a)). This mask was mounted within a 19 mm diameter aperture of an opaque metal disk that fit onto a rotation stage (Newport URS100BPP) with the center of rotation approximately at the center of the aperture.

A very low-power copper anode x-ray tube was used to generate x-rays (TruFocus TFS 6050 Cu, 50 W maximum). It was operated at an anode voltage of 10.1 kV to reduce high-energy bremsstrahlung. A 50 micron thick nickel filter was used to preferentially remove the K_β of the tube spectrum to produce an approximately monochromatic x-ray beam of 8-keV Cu K_α radiation. The rotation stage and aperture were mounted on the end of a 45 cm flight-path to produce a nearly flat-field illumination of x-rays across the 19 mm sample.

The x-ray mask and a static x-ray image of the pattern are shown in Fig. 1(a) and Fig. 1(b). Cornell's LCLS Pixel Array Detector (PAD), comprising a single-chip 194 x 185 pixel array, was placed after the mask along the flight path. The PAD was positioned so the entirety of the




Fig. 2. (a-c) Three sample frames from the 2.5 photon/frame data set with detected x-ray photons circled. (d) Occupancy histogram compared with the Poisson distribution. (e) The sum of all thresholded frames from the 2.5 photon/frame data set showing a uniform angular distribution of data.

aperture image was incident on the detector. The x-ray image shown in Fig. 1(b) was collected by summing 432 images of the mask at fixed position. Each 0.1 s image had an occupancy of approximately 0.2 x-ray photons per pixel per frame in the unobstructed regions.

Very low-fluence data were also collected with a continuously rotating sample. The detector collected images at 100 frames per second with a per-frame integration time of 100 microseconds. The waveforms used for digitization and detector readout were the same as those used when the detector is running at 120 frames per second, the frame rate of the LCLS, except the internal trigger was set to a 10 ms period (rather than 8 ms). X-ray tube currents of 0.15 mA, and 0.03 mA were used to produce varying x-ray intensities. With each current, hundreds of thousands of images were collected. Figures 2(a-c) show three typical very low-fluence mask images, in each case consisting of only a few x-rays per frame.

Dark signal measurements were made throughout the data collection sequence by periodically taking groups of 144 frames with the x-ray shutter closed. These dark frames were used to define a low-noise zero-level which was subtracted from individual signal frames to extract the x-ray induced signal from the raw detector output.

Analog integrating detectors are required at XFELs because many experiments deliver more than one x-ray photon per pixel per frame (as expected at low scattering angles in single larger particle imaging experiments), and the x-ray pulse is too short for photon counting electronics. Minimum signal threshold values can be applied to the analog data to reject low-level noise [6]. The threshold in this experiment was set to 0.7 x-ray photons (for 2.5 photon/frame data set) or 0.75 x-ray photons (for the 11.5 photon/frame data set). At these thresholds approximately 0.05 and 0.01 false positive photon measurements per frame are expected, using a normal distribution and the previously measured [5] pixel signal-to-noise ratio of 7 for a single 8-keV x-ray. The lower threshold was used for 2.5 photon/frame data because this was the last data set taken and progressively less favorable parameters were chosen to test the robustness of the detector and algorithm. No compensation was used for charge sharing between adjacent pixels, nor were pixel gains individually calibrated. A single, global threshold and nominal pixel gain value were applied across the array.

Separate data sets analyzed included hundreds of thousands of frames with mean fluences of 11.5 photons/frame and 2.5 photons/frame.

4. Results

Reconstructed images are shown in Fig. 1(c) and Fig. 1(d). Figure 1(d) was reconstructed using 450,000 frames of data with an average of 2.5 photons per frame. The reconstruction algorithm used 180 equally spaced 2° steps. Figure 2(e) shows a simple sum of the thresholded frames of data that results in a rotationally smeared image with a uniform angular distribution. Figure 2(d) shows a per-frame occupancy histogram, confirming the expected Poisson distribution. This data set has a total of 1.2 million photons. For comparison, a data set with a similar total number of photons, but a higher per-frame photon average (and thus fewer frames) was also processed. The reconstruction is shown in Fig. 1(c), where the average occupancy was 11.5 photons/frame.

The quality of the two reconstructions differs in both spatial resolution and contrast, with the 11.5 photons/frame data yielding better results. This agrees with the results of reconstructing 3D intensities from simulated single-particle diffraction data [3], also at very low fluence. The degradation in quality occurs when a significant fraction of the information content in each frame, about half, is just the orientational state. There is a sharp increase in the iteration count of the EM algorithm when this criterion is met: the 2.5 photon/frame data required 220 iterations, compared to 49 iterations for the 11.5 photon/frame data. The convergence of these 2D reconstructions is similar to the performance of the algorithm in 3D with simulated data in the ultra-low-fluence limit. This provides further confirmation that our test scenario is directly relevant to single protein imaging at an XFEL. Supplementary materials associated with this paper include a movie ([Media 1](#)) showing progression of the model through 300 iterations from initial guess to converged solution for the 2.5 photon/frame data set.

By adding a uniform distribution of computer generated photon counts to the data sets, and processing it by the EM algorithm as before, we are able to simulate the effects of $SN = 1$ background scattering from gas molecules along the path of the incident x-ray beam in single particle experiments. This should be the major source of background signal and many times larger than the detector noise when the detector data are properly thresholded. Not surprisingly we find deterioration in the quality of the reconstruction. The degree of degradation is consistent with the information ratio R quoted above, which equals 0.26 for our chosen signal-to-noise. With this level of background our data set with 11.5 signal-photons/frame corresponds to a zero-background data set with only 3 photons/frame. The resulting reconstruction by the EM algorithm was therefore similar to that of our 2.5 photons/frame background-free reconstruction in both image quality and number of iterations (Fig. 3). Repetition of the algorithm starting with another initially random guess always results in a reconstruction that is identical to the eye, but with an arbitrary rotation of the final image. This has been verified by repeated trials (data not shown).

5. Conclusion

Although this demonstration was motivated by the ongoing effort to realize single particle imaging, the strategy we employed applies more generally to measurements which seek to eliminate ensemble averaging and as a result yield extremely weak signals. Temporal averaging is avoided by short pulses of illumination and the spatial counterpart is achieved by isolation (e.g. single particles) or focusing, as in the case of ptychography [11]. In all these cases one sacrifices signal within a single frame, thus putting an increased burden on the recording of weak signals with high fidelity and reconstructing from the resulting very sparse data. The envisioned single particle experiments at XFELs are an extreme example of this, but the same approach would




Fig. 3. Effect of background on reconstruction quality. (a) Reconstruction from 2.5 photons/frame data set and no added background. This is the same as Fig. 1(d). (b) Reconstruction from the 11.5 photons/frame data set with an average of 11.5 photons of background added per frame ‘by hand’ with a Poisson distribution. The background level was subtracted before rendering to facilitate comparison to (a). As can be seen, the quality of the reconstructions is about the same, and much reduced from the original 11.5 photons/frame data (Fig. 1(c)).

also apply to experiments with lower intensity sources, for example, Energy Recovery Linac (ERL) x-ray sources [12]. An ERL can deliver very short x-ray pulses that are much less intense than XFEL pulses, but deliver many more pulses per second to compensate. Ptychography performed with an ERL, in conjunction with our data acquisition/analysis method, looks especially promising. Data acquisition would be fast and yet immune to mechanical instabilities because of the short pulse duration, while jitter in the position of the focus would be algorithmically reconstructed, in analogy with the angular reconstructions in our demonstration.

Appendix: Information reduction by background

The effect of background in low-flux x-ray measurements is well modeled by a noisy communications channel [13]. Each detector pixel represents one channel and the noise analysis can be carried out for a single channel because the noise processes are, to a good approximation, already independent at the level of pixels. (Background scattering from gas molecules is incoherent and uncorrelated between pixels, as is the photo-absorption process that in effect samples the number of photons in the radiation field.)

Let w be the x-ray flux integrated over one pixel in the time interval of a single frame. In our experiment w takes two values: the background value v when the pixel is covered by mask, and $v + \mu$ when the pixel is uncovered. Because the mask is given random rotations, the prior distribution on w is,

$$p(w) = (1 - \sigma)\delta(w - v) + \sigma\delta(w - v - \mu), \quad (5)$$

where σ represents the fraction of open area in the mask.

The detector pixel measures w as a discrete number of photons k . This is a Poisson process




Fig. 4. Reduction in the rate at which pixels measuring photons acquire information as a function of signal-to-noise. Plotted is the ratio R of the information rate, acquired with and without background. SN is the ratio of signal to background photon counts. This plot applies to the model where half the pixels receive only background (covered by mask) and the other half receive signal and background (not covered by mask).

with conditional probability

$$p(k|w) = \frac{w^k}{k!} e^{-w}. \quad (6)$$

In the communications analogy k is the message that is received when w is sent. The information obtained in the measurement (transmitted by the channel) equals the mutual information I associated with the joint probability distribution

$$p(w, k) = p(w)p(k|w). \quad (7)$$

The mutual information is the difference of entropies

$$I = H_k - H_{k|w}, \quad (8)$$

where

$$H_k = \sum_k -p(k) \log p(k) \quad (9)$$

is the entropy of the photon counts and

$$H_{k|w} = \int dw p(w) \sum_k -p(k|w) \log p(k|w) \quad (10)$$

is the average entropy of the counts when the flux is given.

Both entropies simplify considerably in the extreme low flux limit ($\nu \rightarrow 0$, $\mu \rightarrow 0$) where we can neglect $k > 1$ in the sums:

$$I(\mu, \nu) = \mu (\sigma(1 + \text{SN}^{-1}) \log(1 + \text{SN}) - (\sigma + \text{SN}^{-1}) \log(1 + \sigma \text{SN})) + O(\mu^2). \quad (11)$$

Here $SN = \mu/v$ is the signal-to-noise. The zero background limit at low signal flux corresponds to the limit of infinite SN in the expression above:

$$I(\mu, 0) = -\mu \sigma \log \sigma + O(\mu^2). \quad (12)$$

The ratio of these quantities depends only on SN and represents the ratio of the information rate is acquired with and without background:

$$\frac{I(\mu, v)}{I(\mu, 0)} = R(SN) = \frac{\sigma(1 + SN^{-1}) \log(1 + SN) - (\sigma + SN^{-1}) \log(1 + \sigma SN)}{-\sigma \log \sigma}. \quad (13)$$

The function R is linear at small SN and monotonically approaches 1 at large SN. Figure 4 shows a plot for the case $\sigma = 1/2$.

Acknowledgments

LCLS PAD development was supported by subcontract from SLAC under DOE Contract DE-AC02-76SF00515. Detector development at Cornell is also supported by DOE Grants FG02-97ER62443, DE-FG02-10ER46693 and the Keck Foundation. CHESS is supported by NSF and NIH-NIGMS under NSF Grant DMR-0936384. The data analysis dealing with the extraction from randomly-oriented, sparse data is supported by DOE Grant DE-FG02-11ER16210.