



This is the accepted manuscript made available via CHORUS. The article has been published as:

Emission ghost imaging: Reconstruction with data augmentation

K. J. Coakley, H. H. Chen-Mayer, B. Ravel, D. Josell, N. N. Klimov, D. S. Hussey, and S. M. Robinson

Phys. Rev. A **109**, 023501 — Published 1 February 2024

DOI: [10.1103/PhysRevA.109.023501](https://doi.org/10.1103/PhysRevA.109.023501)

1 **Emission Ghost Imaging: reconstruction with data augmentation**

2 K.J. Coakley*

3 *National Institute of Standards and Technology*

4 *325 Broadway, Boulder, CO 80305 USA*

5 H.H. Chen-Mayer, B. Ravel, D. Josell, N.N. Klimov, and D.S. Hussey

6 *National Institute of Standards and Technology*

7 *100 Bureau Drive, Gaithersburg, MD 20899 USA*

8 S.M. Robinson[†]

9 *PREP Associate, Physical Measurement Laboratory*

10 *National Institute of Standards and Technology*

11 *Gaithersburg, MD 20899 USA and*

12 *Department of Materials Science and Engineering*

13 *University of Maryland, College Park, MD 20742-2115 USA*

14 (Dated: January 5, 2024)

Abstract

Ghost Imaging enables 2D reconstruction of an object even though particles transmitted or emitted by the object of interest are detected with a single pixel detector without spatial resolution. This is possible because for the particular implementation of ghost imaging presented here, the incident beam is spatially modulated with ~~a attenuating mask with many distinct orientations where each orientation yields a distinct spatial pattern in the attenuated beam.~~ **non-configurable attenuating mask whose orientation is varied (e.g. via transverse displacement or rotation) in the course of the ghost imaging experiment.** ~~many distinct orientations where~~ Each orientation yields a distinct spatial pattern in the attenuated beam. In many cases, ghost imaging reconstructions can be dramatically improved by factoring the measurement matrix which consists of measured attenuated incident radiation for each of many orientations of the mask at each pixel to be reconstructed as the product of an orthonormal matrix Q and an upper triangular matrix R provided that the number of orientations of the mask (N) is greater than or equal to the number of pixels (P) reconstructed. For the $N < P$ case, we present a data augmentation method that enables QR factorization of the measurement matrix. To suppress noise in the reconstruction, we determine the Moore-Penrose pseudoinverse of the measurement matrix with a truncated singular value decomposition approach. Since the resulting reconstruction is still noisy, we denoise it with the Adaptive Weights Smoothing method. In simulation experiments, our method outperforms a modification of an existing alternative orthogonalization method where rows of the measurement matrix are orthogonalized by the Gram-Schmidt method. We apply our ghost imaging methods to experimental X-ray fluorescence data acquired at Brookhaven National Laboratory.

* kevin.coakley@nist.gov

† Current address: Sensor Science Division, Physical Measurement Laboratory, National Institute of Standards and Technology, Gaithersburg, MD 20899 USA

I. INTRODUCTION

Although first applied to quantum optics [1, 2] ghost imaging also has classical applications including: experiments with pseudorandom light [3]; X rays [4–7]; electrons [8]; neutrons [9]; and lidar [10] (all of which are in transmission mode). Recently, an implementation of ghost imaging, in emission mode, was developed for elemental mapping based on X-ray fluorescence signals [11, 12]. Here, we focus on ghost imaging of spatially varying emission yields. Ghost imaging is attractive because it enables ~~enables~~ **may enable** high quality reconstructions with lower doses of incident radiation than doses associated with traditional methods [13]. Further, detected particles need not be collimated in ghost imaging. This aspect of ghost imaging is important because in some applications (~~e.g. prompt gammas gamma rays produced by neutrons~~) **(e.g. possible emission ghost imaging based on prompt gamma rays produced by neutron illumination [9])**, collimation is not feasible.

In ghost imaging **with penetrating radiation such as neutrons or X-rays**, the incident beam is attenuated by a **non-configurable mask whose orientation is varied (e.g. via transverse displacement or rotation) in the course of the ghost imaging experiment**. At each mask orientation, the spatial variation of the attenuated incident beam is distinct. Without the object of interest, one acquires measurements of the attenuated incident beam at each pixel due to each mask orientation with a position sensitive detector. This data forms the measurement matrix A . With only the object of interest placed behind the mask, for each mask orientation, ~~a fraction of the total number of particles transmitted through the object~~ **(or particles emitted by the object) is (ideally) detected by a single pixel detector termed the bucket detector**. **In practice, the total solid angle subtended by the bucket detector about any location in the object where emission occurs is less than 4π . This** Hence, the expected **fractions of emitted particles that are detected in emission ghost imaging experiments** depend on solid angle effects and detector efficiency. Reconstructions are determined given the “bucket” data and the measurement matrix.

In many applications, ghost imaging reconstructions can be dramatically improved when the measurement matrix A is factored as the product of an orthonormal matrix Q and an upper triangular matrix R (see, for example, [6]). However, QR factorization of A is not feasible when the number of mask orientations N is less than the number of pixels P that are reconstructed. As discussed in [14], the $N < P$ case is not rare. For the $N < P$ case,

we present, to the best of our knowledge, a new data augmentation method that increases N and enables QR factorization of the augmented measurement matrix. We remark that noisy training data is often extended in machine learning studies with data augmentation methods (see for example [15, 16]). However, these methods are very different from our data augmentation method.

Our reconstruction algorithm based on the QR factorization of the measurement matrix is similar to the method in [6]. However, unlike in [6], we suppress noise effects by determining the Moore-Penrose pseudoinverse [17–19] of the measurement matrix with a truncated singular value decomposition (TSVD) method [20, 21]. We note that Chen et al. [22] implemented a TSVD method for ghost imaging with a reconstruction method that differs from our QR method. In general, reconstructions obtained with our method are still noisy. Hence, we denoise them with the Adaptive Weight Smoothing (AWS) method [23–25]. Like wavelet methods, AWS can smooth out noise while preserving edges in images. For other denoising methods for ghost imaging, see, for example, [26–28].

In Section II, we discuss our ghost imaging reconstruction method, our data augmentation method, and the Adaptive Weights Smoothing method. In Section III A, in a simulation study, we demonstrate that our method yields reconstructions with lower root-mean-square error (RMSE) than reconstructions obtained by a modified version of an existing method [14] where the rows of the measurement matrix are orthogonalized with a Gram-Schmidt procedure rather than a QR method. In Section III B, we apply our ghost imaging methods to experimental X-ray fluorescence data acquired at Brookhaven National Laboratory. In Section IV, we discuss results and summarize some diagnostic studies.

II. METHODS

A. Reconstruction method: no QR step

First, we define the measurement matrix A . Since the photon imaging detector acquires a two-dimensional (2D) image of size 6 by 46, there are 276 pixels. We assign an index to each pixel. For instance, the pixel indexes in the first row of the photon image detector range from 1 to 46. In the second row, they range from 47 to 92. And so on. The ij (th) element of A corresponds to the measured attenuated incident beam produced by the i th

mask orientation at the j th pixel. The size of A is N by P where N is the number of mask orientations and P is the number of pixels. For instance, if we reconstruct an image of size 6 by 46, $P = 276$. We define the theoretical measurement matrix to be $\tilde{A}$. By theoretical, we mean what would be observed in an ideal experiment without measurement errors. We define the observed and theoretical bucket data vectors to be b and $\tilde{b}$ respectively. Both b and $\tilde{b}$ are N -dimensional vectors. The i th component of b , b_i , is the observed number of events detected by the bucket detector for the i th orientation of the attenuating mask. Given the theoretical measurement matrix $\tilde{A}$ and theoretical emission yield vector $\tilde{v}$, we have that

$$\tilde{A}\tilde{v} = \tilde{b}. \quad (1)$$

This theoretical relationship applies to transmission ghost imaging as well. For experimental data, due to ~~random and possible systematic errors~~ **measurement error**, we have that

$$Av \approx b. \quad (2)$$

Based on A and b , our estimate of $\tilde{v}$ is v . Below, we describe our approach to estimate $\tilde{v}$.

The ghost imaging estimate of $\tilde{v}$ (which can be an attenuation term for transmission studies or an emission term like in our study) at the pixel centered at (x, y) , $v(x, y)$, satisfies the following equation (see Eq. 7 of [6])

$$v(x, y) * PSF(x, y) = \frac{1}{N} \sum_{i=1}^N (b_i - \bar{b}) I_i(x, y), \quad (3)$$

where $*$ denotes convolution, where N is the number of mask orientations, and $I_i(x, y)$ is the measured attenuated incident beam at location (x, y) due to the i th orientation of the attenuating mask, and $\bar{b} = \frac{1}{N} \sum_{i=1}^N b_i$. In the ghost imaging literature, for masks that are spatially random, the set of $I_i(x, y)$ is typically referred to as a speckle basis. To simplify calculations, we assume that $(x, y) = \vec{x}_j$ where j is an integer between 1 and P and $\vec{x}_j = (x_j, y_j)$ is the midpoint of the j th pixel. Similarly, $\vec{x}_k = (x_k, y_k)$ is the midpoint of the k th pixel. For estimation of $\tilde{v}$, based on the “approximate completeness relation” discussed in Section 4.1 of [6] and Eq. 14 in [6], we model the PSF as

$$PSF(\vec{x}_j - \vec{x}_k) = \frac{1}{N} \sum_{i=1}^N (I_i(\vec{x}_k) - \bar{I}(\vec{x}_k)) (I_i(\vec{x}_j) - \bar{I}(\vec{x}_j)) \delta_{j,k}, \quad (4)$$

98 where

$$\bar{I}(\vec{x}_j) = \frac{1}{N} \sum_{i=1}^N I_i(\vec{x}_j),$$

99

$$\bar{I}(\vec{x}_k) = \frac{1}{N} \sum_{i=1}^N I_i(\vec{x}_k), \quad (5)$$

100 and $\delta_{j,k}$ is the Kronecker delta function. We get that

$$v(\vec{x}_j) * PSF(\vec{x}_j) = \sum_{k=1}^P v(\vec{x}_k) PSF(\vec{x}_j - \vec{x}_k) =$$

101

$$\frac{v(\vec{x}_j)}{N} \sum_{i=1}^N (I_i(\vec{x}_j) - \bar{I}(\vec{x}_j))^2. \quad (6)$$

102 Rewriting with $\vec{x}_j$ substituted with (x, y) yields

$$v(x, y) * PSF(x, y) = \frac{v(x, y)}{N} \sum_{i=1}^N (I_i(x, y) - \bar{I}(x, y))^2, \quad (7)$$

103 where

$$\bar{I}(x, y) = \frac{1}{N} \sum_{i=1}^N I_i(x, y). \quad (8)$$

104 Further simplifying, we get that

$$v(x, y) = \frac{1}{N-1} \sum_{i=1}^N (b_i - \bar{b}) \frac{I_i(x, y)}{\hat{\sigma}_{I(x, y)}^2}, \quad (9)$$

105 where $\hat{\sigma}_{I(x, y)}^2$ is the sample variance of $I_1(x, y), I_2(x, y), \dots, I_N(x, y)$. For a definition of the
 106 sample variance, see, for example, [29]. We note that our Eq. 9 result has some similarity
 107 with a result based on random matrix theory (see Eq. 14 of [7]).

108 In our simulation studies, each photon is detected with probability 1. For cases where
 109 each photon is detected with probability $p_{det} < 1$, ~~one would rescale the reconstruction at~~
 110 ~~any pixel by $1/p_{det}$~~ one would scale $v(x, y)$ by $1/p_{det}$. A similar comment may apply to
 111 emission ghost imaging of experimental data.

112 Based on Eq. 9, we can express v as

$$v = \frac{1}{N-1} W A^T (b - \bar{b} \vec{1}), \quad (10)$$

113 where $\vec{1}$ is a N -dimensional vector where each component is 1, A^T is the transpose of A , and
 114 W is a diagonal matrix of size P by P where the k th diagonal value of W is $1/\hat{\sigma}^2(k)$. We
 115 note that in many ghost imaging papers, W is neglected. After we obtain the N -dimensional
 116 v , we form the 2-D reconstruction of interest by mapping the components of the vector v
 117 into a matrix.

118 B. Reconstruction with ~~QR~~ QR method

119 Following Appendix B of [6], for the case $N > P$, we factor the measurement matrix A
 120 as

$$A = QR, \quad (11)$$

121 where Q is an orthonormal matrix of size N by P and R is an upper triangular matrix
 122 of size P by P . We replace A with Q , and b with QA^+b , where A^+ is the Moore-Penrose
 123 pseudoinverse of A . To suppress noise effects, we determine A^+ based on the TSVD of A . In
 124 this approach, terms in the singular value decomposition (SVD) with singular values below
 125 the product of the maximum singular value and an adjustable relative threshold, κ_{SVD} , are
 126 excluded. For detailed discussions on computation of the truncated SVD of A and A^+ (based
 127 on the truncated SVD of A) see [20, 21]. We obtain the reconstruction of $\tilde{v}$ with Eq. 4 (or
 128 equivalently Eq. 5). Similar to methods in [6], we obtain a reconstruction for each of many
 129 permutations of the columns of A (before the QR step). For each pixel, the final recon-
 130 struction is the median value of all the reconstructions. Since the reconstruction depends
 131 on the order of the columns, the permutation method can suppress noise effects. To get a
 132 permutation of the columns of A , we simulate a permutation of the integers $(1, 2, 3, \dots, P)$.
 133 (As in illustration, for the simple case where $P = 4$ and the simulated values of the per-
 134 muted integers are $(3, 1, 4, 2)$, the first, second, third and fourth columns of the new matrix
 135 A would be the third column of the original A matrix, the first column of the original A
 136 matrix, the fourth column of the original A matrix and the second column of the original
 137 A matrix respectively. Because the columns are permuted, the components of v are too.
 138 That is, the first, second, third and fourth components of v correspond to estimates for pixels
 139 3, 1, 4, 2. Hence, the components are re-ordered so that the order of the reconstructed values
 140 correspond to the first, second, third and fourth pixel.) For simulated and experimental data,

the number of permutations are 51 and 101 respectively. In our simulation study, for each value of N , we determine the RMSE of each reconstruction on a grid in κ_{SVD} -space by Monte Carlo simulation. For each N , we select the value of κ_{SVD} that yields the lowest RMSE. For experimental data, we determine reconstructions at various values of κ_{SVD} and select the best reconstruction by scientific judgement. For example, for reconstructions of experimental data, we vary κ_{SVD} over a grid and select the value of κ_{SVD} that appears to produce the reconstruction that best resembles the physical phantom. As a caveat, as discussed in Section IV, in future research, we plan to develop data-driven methods to select κ_{SVD} .

C. Data Augmentation

For the $N < P$ case, we augment A and the bucket signal b . In our simulation studies $P = 276$ and N varies from 50 to 275. For each column of A , $c = (c_1, c_2, \dots, c_N)^T$, we stack k copies of c to form an augmented column c_{aug} . The stacking parameter k is increased until $kN \geq P$. For instance, if $k = 2$, the augmented column is $c_{aug} = (c_1, c_2, \dots, c_N, c_1, c_2, \dots, c_N)^T$. For the case $k = 3$, the augmented column $c_{aug} = (c_1, c_2, \dots, c_N, c_1, c_2, \dots, c_N, c_1, c_2, \dots, c_N)^T$. Extensions for larger k are similar to the above. In our augmented measurement matrix, each column has a periodic structure. The bucket signal is augmented in a similar manner, and also has a periodic structure. The dimension of the augmented bucket data vector is the same as the number of rows in the augmented measurement matrix.

D. Reconstruction with row-orthogonalization method

For the $N < P$ case, Luo et al. [14] orthogonalized the rows of the measurement matrix with a Gram-Schmidt method. Given the projection coefficients determined in the Gram-Schmidt method, the bucket signal was also transformed (see [14] for more details). Based on the transformed versions of A and b , the reconstruction is

$$v = \frac{\alpha}{N} A^T (b - \bar{b} \vec{1}), \quad (12)$$

where α is an adjustable scale parameter. Initially, we determine v with Eq. 9 with $\alpha = 1$. Based on the predicted bucket signal, $\hat{b} = Av$, and the observed bucket signal b , we estimate

the scaling factor α as $\hat{\alpha}$ where,

$$\hat{\alpha} = \frac{b_1 + b_2 \cdots + b_N}{\hat{b}_1 + \hat{b}_2 \cdots + \hat{b}_N}. \quad (13)$$

We then adjust v by scaling it by $\hat{\alpha}$. We stress that this method for determination of α is not discussed in [14]. In our primary simulation study, for the cases studied, $\hat{\alpha}$ ranges from approximately 60 to 300.

E. Denoising with Adaptive Weights Smoothing

We denoise reconstructions with the Adaptive Weight Smoothing (AWS) method [23–25]. In this approach, local polynomial models are fit to 2D reconstructions by maximizing a weighted likelihood function in the neighborhood about each spatial location. At each spatial location, the denoised value is the value predicted by the local polynomial model. In the AWS method, the weighted likelihood function at point x is modeled as

$$L(W(x), \theta) = \sum_i w_i(x) \log p(Y_i, \theta), \quad (14)$$

where Y_i is measured at x_i , w_i is the weight corresponding to the measured value at x_i and the point of interest at x , p is the likelihood of Y_i given the polynomial model parameter vector θ . The basic idea of the AWS approach is to adaptively select the size of the neighborhood about any point of interest and the associated weights in that neighborhood. The overall smoothness of the resulting image depends on the choice of a bandwidth parameter, h_{max} , that specifies the maximum size of any local neighborhood, and a parameter λ that determines when to stop expanding the size of the local neighborhood about any point according to a hypothesis test criterion. In general, as λ increases, the resulting image becomes smoother. Like wavelet methods, the AWS method can preserve edges while smoothing within regions between jumps. In our studies, we equate h_{max} to 2.5 pixel lengths, $\lambda = 1$ and the degree of the local polynomial is 2. We implement AWS in the statistical computing language and environment R [30] by calling the function **lpaws** in the R package **aws**.

190 III. RESULTS

191 A. Simulation Study

192 Our simulation model is based on experimental data acquired at Brookhaven National
193 Laboratory (BNL). In the experiment, synchrotron X rays produce fluorescent ~~photons~~ **X-**
194 **rays** from a physical phantom. Since the incident beam flux was too narrow in the vertical
195 direction (1 mm) to illuminate the entire physical phantom (which is 2.5 mm ~~wide~~ **tall** in
196 the vertical direction), the physical phantom was translated vertically three times to ensure
197 that the bucket data for all three regions of the physical phantom had sufficiently high
198 signal-to-noise. We simulate one measurement matrix that applies to all three translations
199 of the physical phantom. For each translation, we simulate distinct bucket data and obtain
200 a reconstruction. We patch the reconstructions together to form an overall reconstruction
201 (see section III B for more experimental details).

202 In FIG. 1, we show the measurement matrix for the BNL experiment. The average
203 number of counts per element of the measurement matrix is 2.56×10^5 . The average number
204 of counts per component of the bucket data vectors is 3.64×10^6 .

205 In our primary simulation study, the theoretical (true) measurement matrix, $\tilde{A}$, is equated
206 to a scaled version of observed measurement matrix shown in shown in FIG. 1 where the
207 scaling factor is 10. (When a matrix is multiplied by a scalar, each element of the matrix
208 is multiplied by the scalar.) The digital phantom, $\tilde{v}$, is scaled so that the sum of the total
209 expected counts in the three simulated bucket data vectors agrees with 10 times the sum
210 of the counts in the three experimental bucket data vectors. Since the digital phantom (see
211 FIG. 11a) is a rendition of the optical image of the physical phantom (see FIG. 11b), it has
212 complexity similar to the physical phantom. We simulate A by adding Poisson noise to each
213 element of the theoretical measurement matrix $\tilde{A}$. For each translation, we simulate bucket
214 data by adding Poisson noise to the theoretical bucket signal $\tilde{b}$ (see Eq. 6). We simulate A
215 (with size 442 by 276) and b with dimension 442, and then select subsets of A and b – the
216 first N rows of A and the first N components of b .

217 In FIG. 2, we show reconstructions of noise-free data for $P = 276$ where N varies from
218 100 to 275, and scatterplots of reconstructions and the digital phantom. We determine these
219 reconstructions with our data augmentation method with the QR and TSVD methods. Since

the data are noise-free, we set the relative threshold for the truncated SVD, κ_{SVD} , to the very low value of 1.49×10^{-8} based on machine precision considerations. Here, we determine the TSVD with the function **ginv** in the R package **MASS** [31].

In our primary simulation study, we simulate 100 realizations of A and b for each of many values of N . For each realization of A and b , we determined a reconstruction for each of many values of κ_{SVD} on a grid. For each value of N and κ_{SVD} , we estimate the root-mean-square-error (RMSE) of the reconstruction. For each N , $\kappa_{SVD} = 10^{-4}$ yields the lowest RMSE. As a caveat, for $N = 50, 75, 100$ values of κ_{SVD} less than 0.0001 also yield the minimum value. If multiple values of κ_{SVD} yield the minimum RMSE, we report the largest of these values. In FIG. 3, we show reconstructions of noisy simulated data for the $N = 275$ case for various values of κ_{SVD} . In FIG. 4, we show how Monte Carlo estimates of RMSE vary with κ_{SVD} for various values of N . In FIG. 5, we show reconstructions of simulated noisy data and associated scatterplots. In FIG. 6, we show denoised versions of the FIG. 5 reconstructions and associated scatterplots. In FIG. 7, we compare denoised reconstructions determined with our method and the method from [14].

For all cases considered, the RMSE of reconstructions obtained with our method are lower than the RMSE of reconstructions obtained with the method from [14] (see Table I). For our method, RMSE decreases as N increases (see Table I). For the method from [14], RMSE decreases as N increases from 50 to 150, but then increases as N increases from 150 to 275. The authors of [14] observed a similar instability as N increased, and attributed the phenomenon to numerical instability effects. For both methods, denoised reconstructions have lower RMSE than reconstructions that are not denoised with AWS (see Table I).

TABLE I. Estimated expected value of root-mean-square error (RMSE) of reconstructions for primary simulation study. Theoretical measurement matrix equated to 10 times the experimental measurement matrix shown in FIG. 1. (a) estimated RMSE of data augmentation reconstruction. (b) estimated RMSE of reconstruction determined with method from [14]. (c) ratio of b and a. (d) estimated RMSE of denoised data augmentation reconstruction. (e) estimated RMSE of denoised reconstruction determined with method from [14]. (f) ratio of e and d.

N	(a)	(b)	(c)	(d)	(e)	(f)
50	0.0608	0.065	1.07	0.0616	0.0643	1.04
75	0.0570	0.0604	1.06	0.0555	0.0569	1.02
100	0.0547	0.0580	1.06	0.0524	0.0531	1.01
150	0.0499(1)	0.0522(1)	1.04	0.0427	0.0427	1.00
200	0.0470(1)	0.0644(3)	1.37(1)	0.0384	0.0400(1)	1.04
225	0.0452(1)	0.0811(6)	1.79(1)	0.0367	0.0413(1)	1.12
250	0.0435(1)	0.1310(14)	3.01(3)	0.0352(1)	0.0505(2)	1.44(1)
260	0.0429(1)	0.1880(26)	4.39(6)	0.0346	0.0642(5)	1.85(2)
265	0.0428(1)	0.2450(42)	5.71(10)	0.0344(1)	0.0789(9)	2.29(3)
270	0.0425(1)	0.3780(130)	8.88(31)	0.0341(1)	0.1090(17)	3.20(5)
275	0.0423(1)	0.8020(441)	19.00(104)	0.0340(1)	0.3610(348)	10.60(102)

B. Experimental Results

We acquired experimental data at the National Synchrotron Light Source II at Brookhaven National Laboratory (BNL) at the NIST Beamline for Materials Measurement (BMM) (see FIG. 8). The energy of the incident X-ray beam was 12 keV (see FIG. 10). The cross-sectional area of the incident X-ray beam, set by 4-blade tungsten slits, was 8 mm x 1 mm. We acquired measurement matrix data with a Dectris Pilatus 2D silicon hybrid photon counting imaging detector. In this detector, the pixel width is approximately 0.172 mm. We monitored beam intensity with ionization chambers (standard beamline equipment) during both the setup stage and the primary stage of the experiment.

The barcode mask was fabricated by patterning and etching arrays of trenches in a Si wafer, metallizing the entire surface with a sputtered Au seed layer and then applying an

additive-based electrochemical deposition process to selectively fill the trenches with gold from the bottom upward [32, 33]. Distinct orientations of the mask were produced by translating the mask (ideally) in the x-direction (see FIG. 9). For $i = 1, 2, \dots, 442$, the i th translation was $(i-1) 50 \mu\text{m}$. Since there are 276 pixels in the imaging detector, the size of the measurement matrix is 442 by 276 (see FIG. 1).

The physical phantom consists of a copper wire (see FIG. 11b) that is shaped to spell out “NIST”. When placed behind the mask and irradiated with X rays, it emitted a fluorescence light X-ray fluorescence spectrum. At each mask orientation, for each translation of the physical phantom, the associated bucket signal is the number of counts in the Cu k-edge fluorescence peaks. This spectrum was acquired with a 4-element Si drift detector with Quantum Electronic XSpres 3X counting electronics.

As mentioned earlier, the narrow beam width necessitated three measurements to cover the entire height of the phantom. Bucket data was acquired for each of three translations of the physical phantom. The translations (in the y-direction) were: 0δ , 6δ , and 12δ where $\delta \approx 0.172 \text{ mm}$ is the width of each pixel in the imaging detector. The rectangular areas associated with the translations did not overlap. For each translation, we obtain a reconstruction of size 6 by 46. We patch these reconstructions together to form an overall reconstruction of size 18 by 46.

In our experimental study, based on the measurement matrix shown in FIG. 1, we determine reconstructions of experimental data without the QR step and with the QR and TSVD steps (see FIG. 11d and FIG. 11h). We show denoised versions of these reconstructions in FIG. 11f and FIG. 11j. Based on visual inspection, the denoised reconstruction of experimental data determined with the QR and TSVD steps (see FIG. 11j) appears to reveal the “N” and “I” features more clearly than the denoised reconstruction determined without the QR step (see FIG. 11f). In both of these reconstructions, the “N”, “I” and “T” features are visible. However, the “S” feature is not visible in either reconstruction. Overall, it appears that the FIG. 11j reconstruction is better than the FIG. 11f reconstruction.

As a diagnostic check, we simulate data based on the the FIG. 1 measurement matrix. In this simulation, the expected number of counts in the observed measurement matrix and the sum of the expected counts in the three bucket data vectors agree with the BNL experiment. The denoised reconstruction of the simulated data determined with the QR and TSVD steps (see FIG. 11i) is vastly superior to the denoised reconstruction of simulated data determined

without the QR and without the TSVD step (see FIG. 11e) However, for the experimental data, the denoised reconstruction with the QR and with the ~~$TSVD$~~ **TSVD** step is better but not vastly superior to the denoised reconstruction without the QR and without the ~~$TSVD$~~ **TSVD** step. Further, the denoised reconstruction of the simulated data determined with the QR and TSVD steps (see FIG. 11i) is vastly superior to the denoised reconstruction of experimental data determined with the QR and TSVD steps (see FIG. 11j) (The above remarks about denoised reconstructions apply to comparisons of reconstructions as well). A plausible systematic error that might explain these results are relative mask orientation errors in the experiment that acquires the measurement matrix, and the experiment that acquires the bucket data.

In an attempt to suppress relative mask orientation errors in the two experiments, we aggregate the measurement matrix and bucket data. In this scheme, the measurement matrix elements in each column are grouped into blocks of size m . For instance, if $m=3$, the first three elements go into the first block. The next three go into the second block, and so on. Elements in each block are summed. A similar procedure applies to the bucket data. For the $m=3$ choice, the number of rows in A is reduced from 442 to 147. Similarly, the number of components in b is reduced from 442 to 147. In three aggregation schemes, the number of rows in A and the dimension of the bucket data are reduced from 442 to 147, 110, and 55 respectively. After data augmentation, the associated resulting denoised reconstructions for the three aggregation schemes are not dramatically improved relative to the denoised reconstruction computed from the full data with the QR and TSVD steps. Based on visual inspection, the aggregation scheme corresponding to 147 rows in A (see FIG. 11k) appears to produce a slightly better result than the other aggregation schemes.

As discussed in Section II B, when reconstructing experimental data with a QR step, we determine a reconstruction for each of 101 permutations of the columns of the measurement matrix. For each pixel, we report the median value of the associated 101 reconstructions. To understand how results vary with the number of permutations, N_p , for the $N = 442$ case, we determine a reference reconstruction at $N_p = 1001$. For this case, reconstructed values fall in the interval $(-0.049, 0.558)$, and denoised reconstructed values fall in the interval $(-0.041, 0.313)$. We determine a reconstruction from the observed measurement matrix ($N_p = 0$), and reconstructions at $N_p = 5, 11, 25, 51, 101$. At each pixel, we compute the difference Δ between each reconstruction and the reference reconstruction and the difference Δ_{dn}

TABLE II. Stability of reconstructions and denoised reconstructions of experimental data determined with QR and TSVD steps as a function of the number of permutations N_p .

N_p	$\min(\Delta)$	$\max(\Delta)$	$\min(\Delta_{dn})$	$\max(\Delta_{dn})$
0	-1.04×10^{-1}	4.98×10^{-2}	-9.02×10^{-2}	2.99×10^{-2}
5	-5.10×10^{-4}	4.63×10^{-4}	-2.37×10^{-4}	2.13×10^{-4}
11	-1.13×10^{-4}	1.63×10^{-4}	-5.78×10^{-5}	9.40×10^{-5}
25	-1.28×10^{-4}	1.35×10^{-4}	-7.45×10^{-5}	2.31×10^{-4}
51	-7.99×10^{-5}	8.91×10^{-5}	-8.38×10^{-5}	2.56×10^{-5}
101	-5.06×10^{-5}	3.45×10^{-5}	-2.25×10^{-5}	1.36×10^{-5}

between each denoised reconstruction and the denoised reference reconstruction. In Table II, we show the minimum and maximum value of Δ and Δ_{dn} for each value of N_p . Based on Table II, it appears that the choice of $N_p = 11$ (or higher) yields sufficiently stable results. For simulated data, we expect that reconstructions (determined with a QR step) are also sufficiently stable for $N_p = 11$ (or higher).

IV. DISCUSSION

Our research on data augmentation was motivated by the goal of understanding X-ray fluorescence emission rate measurements determined with ghost imaging. In future experiments, we plan to investigate possible mask orientation errors that may have affected our measurements.

For reconstructions obtained from simulated data without the QR step and without the row permutation method, it is straight-forward to show that the reconstruction of the original data and the reconstructions of its augmented versions agree exactly (if computed with an infinite precision computer). **This observation is consistent with the claim that data augmentation is a reasonable procedure.**

For the cases studied, for reconstructions determined with the QR step, the predicted number of total counts in the bucket data agrees with the observed number of total counts in the bucket data to six significant digits or more. **This observation is consistent with the claim that our Eq. 9 approach is valid.**

For simulated data shown in FIG. 2e, reconstruction and phantom values are nearly

the same. This observation is consistent with the claim that our Eq. 9 approach is valid. Results shown in FIG. 2f, FIG. 5e and FIG. 5f also support the above claim.

For the case where $N = 442$ and $P=276$, we reconstructed high signal-to-noise simulated data (with $\kappa_{SVD} = 10^{-6}$) as well as its augmented versions where $N = 884, 1326$, and 1768 with QR but without permutation of the columns of the measurement matrix. The RMSE values for reconstructions for $N=442, 884, 1326$ and 1768 are $0.0278, 0.0281, 0.0280$ and 0.0279 . The root-mean-square deviation between the reconstruction of the $N=442$ data, and reconstructions at $N=884, 1326$ and 1768 are $0.0010, 0.0008$ and 0.0011 . Scatterplots show that the reconstructions of the augmented data and the $N = 442$ data are almost the same. Even though data augmentation is not expected to improve reconstructions for the $N > P$ case considered in this simulation study, data augmentation did not significantly degrade the reconstruction. This diagnostic study suggests that data augmentation is reasonable.

In simulation studies, for the $N < P$ case, our method yields yielded reconstructions with lower RMSE compared to the method from [14]. This is plausible because the Gram-Schmidt orthogonalization scheme implemented in [14] may be unstable compared to the QR scheme that we implement in our method. See [7] for discussion of this point. Second, the TSVD step in our method improves results (especially for large N) (see FIG. 4). Another possibility is that our method performs better because we account for the PSF function whereas the method in [14] does not.

In our primary simulation study, the theoretical measurement matrix was equated to 10 times the experimental measurement matrix shown in FIG. 1. For each N , RMSE took its minimum value when $\kappa_{SVD} = 10^{-4}$. In a secondary simulation study, the theoretical measurement matrix was equated to the experimental measurement matrix shown in FIG. 1. For each N , the relative RMSE for the method from [14] increased relative to our method. For $N \leq 100$, RMSE takes its minimum value at $\kappa_{SVD} = 10^{-3}$. For each N above 100, RMSE takes its minimum value at $\kappa_{SVD} = 5 \times 10^{-4}$. Since the data is noisier in the secondary study, it makes sense that a larger κ_{SVD} is required to minimize RMSE.

We remark that compressive sensing has been broadly applied to reflection ghost imaging [34], transmission ghost imaging [10, 35] and recently to emission ghost imaging [11]. Compressive sensing methods are appropriate for cases where signals-of-interest are sparse. In ghost imaging applications, signals-of-interest are typically sparse. For the case where signals-of-interest are not sparse, we expect the performance of compressive sensing methods

to deteriorate. In contrast, we expect our method to apply to both sparse and non-sparse signals-of-interest. We expect compressive sensing to apply to the $N < P$ case. For the data analyzed here, comparison of our method to the compressive sensing method is a worthy topic for future research but beyond the scope of this study. Other methods relevant to the $N < P$ case include pseudoinverse methods [9, 36], singular value decomposition methods [37] and deep learning [27, 38–40].

For analysis of experimental data, we selected the threshold for the truncated SVD method by visual inspection of the reconstructions determined with different thresholds for the truncated SVD method. Future research will focus on data-driven threshold selection methods such as cross-validation or related methods (see, for example, [41]).

As stated earlier, the AWS denoising method is designed to smooth out noise and preserve edges in images. In our simulation study, the denoised versions of the FIG. 5. reconstructions (shown in FIG. 6) have well-preserved edges and noise is significantly suppressed. A similar comment applies to the denoised version of the FIG. 11g reconstruction that is shown in FIG. 11i.

In this work, we reported RMSE statistics. As discussed in [42], there are other candidate metrics as well. For example, how well the mean value of $\tilde{v}$ is determined in a region-of-interest is a possible alternative metric. We note that other metrics were studied in [14]. How well our method performs for higher dimensional cases is a topic for further study.

In our simulation studies, for the cases studied, we demonstrated that denoising ghost imaging reconstructions with AWS reduces RMSE. However, in some analyses, it may be best to extract quantitative information from reconstructions rather than denoised reconstructions. For instance, it may be best to determine the mean emission yield in a region-of-interest as the sample mean of reconstructed emission yields.

We expect that our methods for ghost imaging to be directly applicable to other emission signals including neutron capture prompt gamma rays. Our long-term goal is to develop methods for elemental mapping based on ghost imaging of neutron capture prompt gamma rays.

ACKNOWLEDGMENTS

Research performed, in part, at the NIST Center for Nanoscale Science and Technology. We thank C. Daugherty, M. Frey and J. Prothero of NIST for helpful comments. Certain commercial equipment, instruments, or materials are identified in this paper to foster understanding. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

-
- [1] T. B. Pittman, Y. H. Shih, D. V. Strekalov, and A. V. Sergienko, *Physical Review A* **52**, R3429 (1995).
 - [2] M. J. Padgett and R. W. Boyd, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **375**, 20160233 (2017).
 - [3] Y. Bromberg, O. Katz, and Y. Silberberg, *Physical Review A* **79**, 053840 (2009).
 - [4] H. Yu, R. Lu, S. Han, H. Xie, G. Du, T. Xiao, and D. Zhu, *Physical Review Letters* **117**, 113901 (2016).
 - [5] D. Pelliccia, A. Rack, M. Scheel, V. Cantelli, and D. M. Paganin, *Physical Review Letters* **117**, 113902 (2016).
 - [6] D. Pelliccia, M. P. Olbinado, A. Rack, A. M. Kingston, G. R. Myers, and D. M. Paganin, *IUCrJ* **5**, 428 (2018).
 - [7] D. Ceddia and D. M. Paganin, *Physical Review A* **97**, 062119 (2018).
 - [8] S. Li, F. Cropp, K. Kabra, T. J. Lane, G. Wetzstein, P. Musumeci, and D. Ratner, *Physical Review Letters* **121**, 114801 (2018).
 - [9] A. M. Kingston, G. R. Myers, D. Pelliccia, F. Salvemini, J. J. Bevitt, U. Garbe, and D. M. Paganin, *Physical Review A* **101**, 053844 (2020).
 - [10] C. Zhao, W. Gong, M. Chen, E. Li, H. Wang, W. Xu, and S. Han, *Applied Physics Letters* **101** (2012).
 - [11] Y. Klein, O. Sefi, H. Schwartz, and S. Shwartz, *Optica* **9**, 63 (2022).
 - [12] M. Manni, A. Ben-Yehuda, Y. Klein, B. Lukic, A. Kingston, A. Rack, S. Shwartz, and N. Viganò, *arXiv preprint arXiv:2306.16258* (2023).

- [13] A. M. Kingston, W. K. Fullagar, G. R. Myers, D. Adams, D. Pelliccia, and D. M. Paganin, Physical Review A **103**, 033503 (2021).
- [14] B. Luo, P. Yin, L. Yin, G. Wu, and H. Guo, Optics Express **26**, 23093 (2018).
- [15] C. Shorten and T. M. Khoshgoftaar, Journal of Big Data **6**, 1 (2019).
- [16] T. Hu, T. Tang, and M. Chen, IEEE Access **7**, 125133 (2019).
- [17] E. Moore, Bulletin of the American Mathematical Society **26**, 294 (1920).
- [18] R. Penrose, in Mathematical proceedings of the Cambridge philosophical society, Vol. 51 (Cambridge University Press, 1955) pp. 406–413.
- [19] G. Strang, Linear Algebra and its Applications (Academic Press, New York, 1980).
- [20] P. C. Hansen, T. Sekii, and H. Shibahashi, SIAM Journal on Scientific and Statistical Computing **13**, 1142 (1992).
- [21] N. G. Gençer and S. J. Williamson, IEEE Transactions on Biomedical Engineering **45**, 827 (1998).
- [22] L.-Y. Chen, C. Wang, X.-Y. Xiao, C. Ren, D.-J. Zhang, Z. Li, and D.-Z. Cao, Optics Express **30**, 6248 (2022).
- [23] J. Polzehl and V. Spokoiny, The Annals of Statistics **31**, 30 (2003).
- [24] J. Polzehl and V. Spokoiny, Probability Theory and Related Fields **135**, 335 (2006).
- [25] J. Polzehl, K. Papafitsoros, and K. Tabelow, Journal of Statistical Software **95**, 1 (2020).
- [26] J. Kim, J. Hwang, J. Kim, K. Ko, E. Ko, and G. Cho, Optics Express **29**, 39323 (2021).
- [27] H. Wu, R. Wang, G. Zhao, H. Xiao, J. Liang, D. Wang, X. Tian, L. Cheng, and X. Zhang, Optics and Lasers in Engineering **134**, 106183 (2020).
- [28] X.-R. Yao, W.-K. Yu, X.-F. Liu, L.-Z. Li, M.-F. Li, L.-A. Wu, and G.-J. Zhai, Optics Express **22**, 24268 (2014).
- [29] J. Shao, Mathematical Statistics (Springer Science & Business Media, 2003).
- [30] R Core Team, R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria (2023).
- [31] W. N. Venables and B. D. Ripley, Modern applied statistics with S-PLUS (Springer Science & Business Media, 2013).
- [32] D. Josell, W. A. Osborn, M. E. Williams, and H. Miao, Journal of the Electrochemical Society **169**, 032509 (2022).
- [33] D. Josell, Z. Shi, K. Jefimovs, V. A. Guzenko, C. Beauchamp, L. Peer, M. Polikarpov, and

456 T. P. Moffat, Journal of the Electrochemical Society **168**, 082508 (2021).

457 [34] O. Katz, Y. Bromberg, and Y. Silberberg, Applied Physics Letters **95** (2009).

458 [35] V. Katkovnik and J. Astola, Journal of the Optical Society of America A **29**, 1556 (2012).

459 [36] C. Zhang, S. Guo, J. Cao, J. Guan, and F. Gao, Optics Express **22**, 30063 (2014).

460 [37] X. Zhang, X. Meng, X. Yang, Y. Wang, Y. Yin, X. Li, X. Peng, W. He, G. Dong, and H. Chen,
461 Optics Express **26**, 12948 (2018).

462 [38] Y. He, G. Wang, G. Dong, S. Zhu, H. Chen, A. Zhang, and Z. Xu, Scientific Reports **8**, 6469
463 (2018).

464 [39] L. Zhang, Z. Bian, H. Ye, D. Zhang, and K. Wang, Optics Communications **504**, 127479
465 (2022).

466 [40] S. Rizvi, J. Cao, K. Zhang, and Q. Hao, Scientific Reports **10**, 11400 (2020).

467 [41] A. B. Owen and P. O. Perry, The Annals of Applied Statistics , 564 (2009).

468 [42] H. H. Barrett, J. Denny, R. F. Wagner, and K. J. Myers, Journal of the Optical Society of
469 America A **12**, 834 (1995).




FIG. 1. (a) Experimental measurement matrix acquired at Brookhaven National Laboratory. (b) Scan of measurement matrix (corresponding to red dashed line in (a)).




FIG. 2. The number of pixels is $P = 276$. We equate the theoretical measurement matrix to 10 times the experimental measurement matrix shown in FIG. 1. Here we show reconstructions of noise-free data determined with data augmentation, QR and TSVD steps where the relative threshold for the truncated SVD method, κ_{SVD} , is 1.49×10^{-8} . (a)-(d) Reconstructions for $N=275, 250, 200$ and 100. (e)-(h) Associated scatterplots of reconstructions and the digital phantom. Points that fall on the red line correspond to cases where the reconstruction and true value agree exactly.




FIG. 3. We equate the theoretical measurement matrix to 10 times the experimental measurement matrix shown in FIG. 1. The number of pixels is $P = 276$. We show reconstructions of simulated noisy data where $N = 275$ where κ_{SVD} varies. (a) Digital phantom. (b) to (i) Reconstructions where $\kappa_{SVD} = 10^{-6}, 10^{-5}, 5 \times 10^{-5}, 10^{-4}, 5 \times 10^{-4}, 10^{-3}$ and 5×10^{-3} respectively.




FIG. 4. We equate the theoretical measurement matrix to 10 times the experimental measurement matrix shown in FIG. 1. The number of pixels is $P = 276$. For our data augmentation method (with QR and TSVD steps), we show RMSE of reconstructions of simulated noisy data as a function of κ_{SVD} . (a) $N = 275$. (b) $N = 250$. (c) $N = 200$. (d) $N = 150$.




FIG. 5. We equate the theoretical measurement matrix to 10 times the experimental measurement matrix shown in FIG. 1. The number of pixels is $P = 276$. (a)-(d) Reconstructions of simulated noisy data for $N=275, 250, 200$ and 100. (e)-(h) Associated scatterplots of reconstruction and the digital phantom. Points that fall on the red line correspond to cases where the reconstruction and true value agree exactly.




FIG. 6. Here, we show denoised versions of the FIG. 5 reconstructions. (a)-(d) Denoised reconstructions for $N=275, 250, 200$ and 100 . (e)-(h) Associated scatterplots of denoised reconstruction and digital phantom. Points that fall on the red line correspond to cases where the reconstruction and true value agree exactly.




FIG. 7. We equate the theoretical measurement matrix equated to 10 times the experimental measurement matrix shown in FIG. 1. The number of pixels is $P = 276$. We show denoised reconstructions of simulated noisy data. (a),(c) and (e) Reconstructions determined with the method from [14] for $N = 275, 250$ and 200 . (b),(d) and (f) Reconstructions determined with our method for $N = 275, 250$ and 200 .




FIG. 8. Experimental setup at the NIST Beamline for Materials Measurement at the National Synchrotron Light Source II at Brookhaven National Laboratory.




FIG. 9. The computer-designed barcode mask was fabricated by bottom-up gold filling of nominally $20.7\ \mu\text{m}$ deep trenches patterned and etched into the surface of a silicon wafer. The “gold” bars in the barcode pattern are actually arrays of $8\ \mu\text{m}$ wide Au-filled trenches and $2\ \mu\text{m}$ wide Si spacers that have been repeated as needed for the bar width. Illumination of the system was defined by stepping the barcode mask through a sequence of $50\ \mu\text{m}$ displacements parallel to the barcode, generating a total of 442 spatial patterns that were captured by a CCD camera. The same sequence of illumination patterns was repeated with the specimen in place for X-ray fluorescence data acquisition. The patterns and the corresponding XRF signals were processed offline for ghost imaging reconstruction.




FIG. 10. An X-ray fluorescence spectrum excited by a monochromatic X-ray beam at 12 keV incident on the physical phantom made with copper wire (see photograph in FIG. 11b). The bucket signal is obtained by summing the counts under the Cu K_α and Cu K_β peaks. X rays that undergo elastic scattering or Compton scattering do not contribute counts to the bucket data.






Method \ Phantom		Simulation	Experiment
		(a) 	(b) 
Reconstruction ($N = 442$, $P = 276$) w/o QR	w/o AWS	(c) 	(d) 
	with AWS	(e) 	(f) 
Reconstruction ($N = 442$, $P = 276$) with QR, TSVD	w/o AWS	(g) 	(h) 
	with AWS	(i) 	(j) 
Reconstruction ($N = 147$, $P = 276$) with QR, TSVD	with AWS	(k) 	

FIG. 11. In the “Simulation” column, we show the digital phantom and reconstructions and denoised reconstructions of simulated data with signal-to-noise similar to experimental data. In the “Experiment” column, we show the physical phantom and reconstructions and denoised reconstructions of experimental data. For all cases, the number of pixels is $P = 276$. Except for (k), the number of mask orientations is $N = 442$. In an effort to suppress systematic errors, we aggregate the experimental data so that N is reduced from 442 to 147. To enable analysis with the QR and TSVD steps, we augment the aggregated data so that $N = 294$. In (k), we show the denoised reconstruction of the augmented data.