

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Optically coherent image reconstruction in the presence of phase errors using advanced-prior models

Casey J. Pellizzari, Mark F. Spencer, Charles A. Bouman

Casey J. Pellizzari, Mark F. Spencer, Charles A. Bouman, "Optically coherent image reconstruction in the presence of phase errors using advanced-prior models," Proc. SPIE 10650, Long-Range Imaging III, 106500B (11 May 2018); doi: 10.1117/12.2304692

SPIE.

Event: SPIE Defense + Security, 2018, Orlando, Florida, United States

Optically coherent image reconstruction in the presence of phase errors using advanced-prior models

Casey J. Pellizzari^a, Mark F. Spencer^b, and Charles A. Bouman^c

^aAir Command and Staff College, Maxwell AFB, Montgomery, AL, 36112

^bAir Force Research Laboratory, Directed Energy Directorate, Kirtland AFB, NM, 87111

^cSchool of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, 47907

ABSTRACT

Recently developed coherent-imaging algorithms using Model-Based Iterative Reconstruction (MBIR) are robust to noise, speckle, and phase errors. These MBIR algorithms produce useful images with less signal which allows imaging distances to be extended. So far, MBIR algorithms have only incorporated simple image models. Complex scenes, on the other hand, require more advanced image models. In this work, we develop an MBIR algorithm for image reconstruction in the presence of phase errors which incorporates advanced image models. The proposed algorithm enables optically-coherent imaging of complex scenes at extended ranges.

Keywords: Long-Range Imaging, Digital Holography, Image Reconstruction, Imaging Through Turbulence

1. INTRODUCTON

Optically-coherent imaging systems, such as synthetic aperture LADAR (SAL)¹⁻⁴ and digital holography (DH),⁵ offer significant improvements in resolution and sensitivity compared to passive and non-coherent systems. This makes coherent imaging ideal for long-range imaging applications. In practice, these systems are sensitive to phase errors caused by index-of-refraction perturbations in the atmosphere or optical systems. Fortunately, we can estimate these phase errors directly from the data and apply a digital correction to reconstruct a focused image.

Conventional techniques for reconstructing images from optically coherent data produce estimates of the complex-valued reflection coefficient, g , which is the complex-valued ratio of the reflected field to the incident field. For surfaces which are rough relative to the illumination wavelength, this estimation process leads to images with high-spatial-frequency variations known as speckle. Recently, we developed model-based iterative reconstruction (MBIR) algorithms for jointly computing the maximum a posteriori (MAP) estimates of the phase errors, ϕ , and the real-valued reflectance, r , from SAL and DH data.^{6,7} The reflectance is given by the real-valued ratio of the reflected power to the incident power. Furthermore, the reflectance can be expressed as $r = E[|g|^2]$, where $E[\cdot]$ indicates the expected value. The reflectance, in general, is smoother and has higher spatial correlation as compared to g . We are accustomed to seeing r in conventional images and it is of greater interest for many imaging applications. Additionally, by reconstructing r , we leverage its higher spatial correlation to better constrain the estimation process and produce more accurate estimates of the phase errors with less data and less signal. We have shown that this approach is robust to high noise, speckle, and strong phase errors.⁶⁻⁸

While the MBIR algorithms in Refs. 6-8 work well for simple binary objects with high contrast, they use relatively simple prior models for the reflectance function. We have observed that for gray-scale objects with decreased contrast, those MBIR algorithms do not perform as well. In Ref. 9, we developed an MBIR approach for coherent imaging that incorporated advanced prior models using a modified Plug-and-Play (P&P) framework. While this algorithm can better reconstruct complex objects, it ignores phase errors. Therefore, we can not use it when phase errors are present.

In this paper, we develop a new algorithm for jointly evaluating the MAP estimates of both the reflectance, r , and the phase errors, ϕ , using advanced prior models. To overcome the non-tractable cost function associated

Further author information: (Send correspondence to C.P.)

C.P.: E-mail: casey.pellizzari@us.af.mil, Telephone: 1 505 123 1234

with the coherent data model, we use the expectation maximization (EM) algorithm. Thus, we refer to our new approach as the joint-P&P with EM (JPEM) algorithm. As an example, we consider the case of correcting DH data corrupted by two-dimensional isoplanatic phase errors; however, the concept can be applied to other coherent imaging modalities. We compare the JPEM algorithm to the MBIR algorithm from 7, and the image sharpening (IS) from 10, over a range of different image types and contrasts. Using a quantitative analysis, we show that the JPEM algorithm produces high quality images with better phase-error estimates when compared to either the MBIR or IS algorithms. In particular, when the scene being imaged has low contrast, the JPEM algorithm works much better than MBIR or IS.

2. CONCEPT AND DATA MODEL

Figure 1 shows an example DH system using an off-axis image-plane recording geometry (IPRG).¹¹ The system images an object with a reflectance function, r , and a corresponding reflection coefficient, g , through phase errors, ϕ , located near the pupil plane. The returning signal passes through a pupil with a transmittance function, a , and is imaged onto a focal-plane detector array where it is mixed with a reference field to form a noisy digital hologram with noise power σ_w^2 .

Figure 2 shows example reflectance and reflection-coefficient functions, r and g , respectively, and provides some insight into the basic image processing steps. We extract the signal of interest, y , from the spectrum of the digital hologram. Conventional image-processing techniques form an image by taking an inverse Fast Fourier Transform (FFT) of the data, y . This produces a speckled image, as shown in Fig. 2 (g).

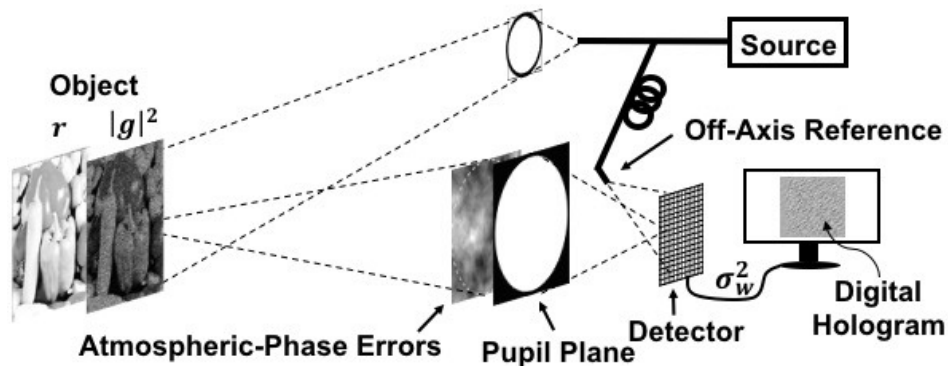


Figure 1. Example DH system using an off-axis IPRG. A laser source is used to flood illuminate an object which has a reflectance function, r , and corresponding reflection coefficient, g . The return signal is corrupted by isoplanatic atmospheric phase errors, ϕ , concentrated near the imaging system. The returning signal passes through the pupil function, a , and is imaged onto a focal-plane detector array where it is mixed with a reference field to form the noisy digital hologram with noise power σ_w^2 .

Using matrix-vector notation, we model the complex-valued pupil-plane image, $y \in \mathbb{C}^M$, obtained from a DH system, as

$$y = A_\phi g + w, \quad (1)$$

where M is the total number of measurement samples produced by the focal-plane array (FPA), $g \in \mathbb{C}^M$ is the object's complex-valued reflection coefficient, and $w \in \mathbb{C}^M$ represents additive Gaussian measurement noise.^{7,8} Finally, the matrix A_ϕ accounts for the propagation and measurement geometry and is a function of the phase errors, ϕ . The matrix is decomposed according to

$$A_\phi = \mathcal{D}(a) \mathcal{D}(\exp\{j\phi\}) F, \quad (2)$$

where $a \in \mathbb{R}^M$ is the pupil-transmittance function, $\phi \in \mathbb{R}^M$ is the phase-error function, and $F \in \mathbb{C}^{M \times M}$ is a two-dimensional discrete Fourier transform (DFT) matrix scaled so that $F^H F = I$.^{*} The superscript, H , indicates the Hermitian transpose.

^{*}For long-range imaging geometries, we ignore the quadratic phase term, Γ , modeled in Ref. 7.

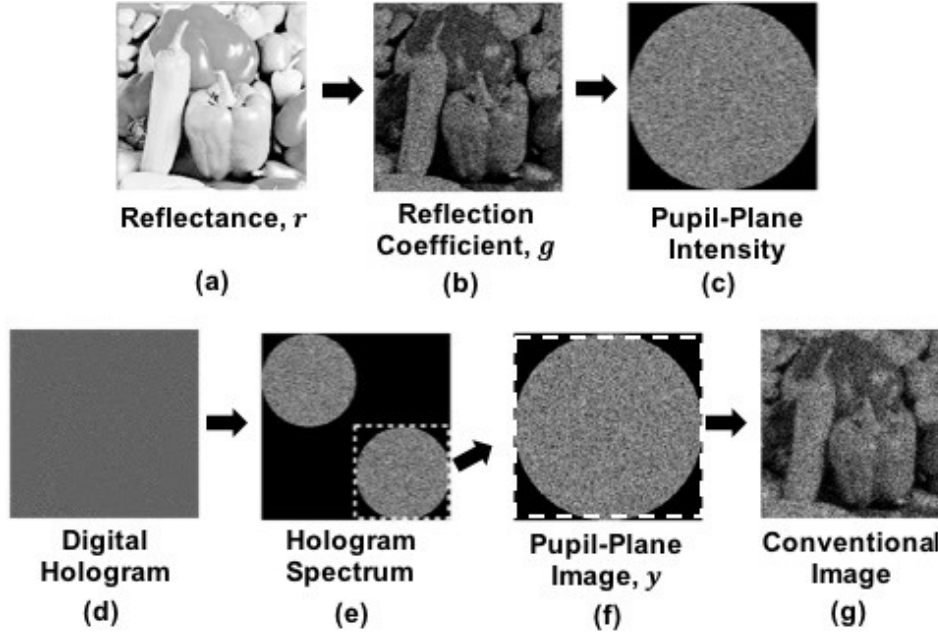


Figure 2. Example quantities for a DH system using an off-axis IPRG. Here, we show the magnitude squared for any complex-valued quantities. An object with the reflectance function shown in (a) has the corresponding reflection coefficient shown in (b) and generates a return signal at the pupil plane of the DH system shown in (c). To extract the signal, we start with a real-valued digital hologram (d), then take a 2D Fast Fourier Transform (FFT) to obtain the complex-valued holographic spectrum (e). Note that the digital hologram, shown in (d), naturally has low contrast. Next, a section of the spectrum is isolated (f), as indicated by the dashed-white line. This subset of the spectrum, y , represents a complex image of the signal field in the pupil plane. Finally, for basic image formation, we take an inverse FFT (IFFT) of y to form the image shown in (g).

3. RECONSTRUCTION FRAMEWORK

For imaging purposes, we wish to estimate the reflectance, $r = E[|g|^2]$, from the pupil image, y , where $|\cdot|^2$ indicates the element-wise magnitude squared. To produce a focused image, we must also estimate and correct the unknown phase errors, ϕ . Therefore, we compute the joint MAP estimates given by

$$\begin{aligned}
 (\hat{r}, \hat{\phi}) &= \underset{r, \phi}{\operatorname{argmin}} \{-\log p(r, \phi|y)\} \\
 &= \underset{r, \phi}{\operatorname{argmin}} \{-\log p(y|r, \phi) - \log p(r) - \log p(\phi)\},
 \end{aligned} \tag{3}$$

where $p(r)$ and $p(\phi)$ are the distributions, or prior models, for r and ϕ . To solve Eq. (3) using an advanced-prior model for $p(r)$, we combine the approaches from 9 and 7. Specifically, we apply the P&P framework to the joint MAP estimation problem of Eq. (3). Then, we introduce a simplified approach using the expectation maximization algorithm.

3.1 Joint Plug and Play

In this section, we apply the P&P framework to compute the joint MAP estimates of both r and ϕ . We start by splitting the variable r into r and x , and rewriting the problem as a constrained optimization given by

$$(\hat{r}, \hat{\phi}) = \underset{r=x, \phi}{\operatorname{argmin}} \{f_1(r, \phi) + \beta f_2(x) + f_3(\phi)\}. \tag{4}$$

Here, $f_1(r, \phi) = -\log p(y|r, \phi)$, $\beta f_2(x) = -\log p(r)$, and $f_3(\phi) = -\log p(\phi)$, where β is a scalar introduced to allow tuning of the prior model for r . We solve Eq. (4) by first forming the unconstrained Augmented Lagrangian

cost function and then we use alternating directions method of multipliers (ADMM).¹² For this problem, the Augmented Lagrangian cost function is

$$L(r, x, \phi; u) = f_1(r, \phi) + \beta f_2(x) + f_3(\psi) + \frac{1}{2\sigma_\lambda^2} \|r - x + u\|^2 - \frac{1}{2\sigma_\lambda^2} \|u\|^2, \quad (5)$$

where u is related to the Lagrange multiplier.¹³ In Eq. (5), the last two terms penalize solutions where $r \neq x$ with a gain set by the parameter σ_λ^2 . The resulting ADMM updates are then given by¹³

$$\begin{aligned} r^{k+1} &= \underset{r}{\operatorname{argmin}} \left\{ f_1(r, \phi^k) + \frac{1}{2\sigma_\lambda^2} \|r - (x^k - u^k)\|^2 \right\} \stackrel{\text{def}}{=} \mathbb{F}(x^k - u^k), \\ x^{k+1} &= \underset{x}{\operatorname{argmin}} \left\{ f_2(x) + \frac{1}{2\beta\sigma_\lambda^2} \|(r^{k+1} + u^k) - x\|^2 \right\} \stackrel{\text{def}}{=} \mathbb{G}(r^{k+1} + u^k), \\ u^{k+1} &= u^k + (r^{k+1} - x^{k+1}), \\ \phi^{k+1} &= \underset{\phi}{\operatorname{argmin}} \{f_1(r^{k+1}, \phi) + f_3(\phi)\}. \end{aligned} \quad (6)$$

In Eq. (6), the inversion operator, \mathbb{F} , is mathematically equivalent to computing the MAP estimate of r using a Gaussian prior model with distribution $p(r) \sim N(x - u, \sigma_\lambda^2 I)$. Here, $N(\mu, C)$ indicates a real-valued normal distribution with mean μ and covariance matrix C .^{14,15} The denoising operator, \mathbb{G} , is mathematically equivalent to computing the MAP estimate of x from the noisy image, $r + u$, corrupted by Gaussian noise. From this perspective, $\sigma_n^2 = \beta\sigma_\lambda^2$ represents the noise variance of an i.i.d. Gaussian forward model and $f_2(x)$ is the prior model.^{14,15} Since \mathbb{G} is equivalent to a Gaussian denoising operation, we use standard Gaussian denoising algorithms (GDAs) to carry out the optimization over x shown in Eq. (6). As a result, we inherit the prior model which is either explicitly or implicitly part of the GDA used. The use of GDAs, and their associated prior models, to solve 6 is the core idea behind the P&P framework.

3.2 JPEM Concept

Unfortunately, for coherent imaging applications, the updates in Eq. (6) can not be evaluated directly due to the high computational complexity of the log likelihood term, $f_1(r, \phi) = -\log p(y|r, \phi)$.^{6,7,9} The exact form of $f_1(r, \phi)$ is found in App. A. As an alternate approach, we simplify the updates using the EM algorithm. Specifically, we introduce the missing data, g , into the log likelihood term to simplify its functional form. Since g is unknown, we remove it from the model by taking a conditional expectation. The resulting simplified function is given by

$$\tilde{f}_1(r, \phi; r^k, \phi^k) = -\mathbb{E} [\log p(y, g|r, \phi) | y, r^k, \phi^k], \quad (7)$$

where $\mathbb{E}[\cdot | y, r^k, \phi^k]$ indicates a conditional expectation taken over g , given the data, y , the current estimate of the reflectance, r^k , and the current estimate of the phase errors, ϕ^k . Equation (7) is known as a surrogate function because it upper bounds $f_1(r, \phi)$ for all r and ϕ . As a result, minimization of \tilde{f}_1 implies minimization of f_1 .¹³ Therefore, we can substitute $\tilde{f}_1(r, \phi; r^k, \phi^k) \rightarrow f_1(r, \phi)$ in Eq. (6) for the updates of x^{k+1} and ϕ^{k+1} .

Our notional concept for using \tilde{f}_1 within the P&P framework is shown in Fig. 3. For each iteration, we construct \tilde{f}_1 once by taking the conditional expectation shown in Eq. (7). This constitutes the E-step of the EM algorithm. Details on the exact form of \tilde{f}_1 can be found in App. B. Next, we apply the ADMM updates from Eq. (6) using the constructed \tilde{f}_1 . This constitutes the M-step of the EM algorithm.

4. JPEM ALGORITHM AND EXECUTION

The notional concept in Fig. 3 consists of three main parts, the inversion operation, the denoising operation, and the phase-error update. In this section, we present details on how to accomplish each of those steps and then present the JPEM algorithm.

Initialize: r^0, ϕ^0, u^0
Repeat{
 E-step:
 $\tilde{f}_1(r, \phi; r^k, \phi^k) = -\text{E} [\log p(y, g|r, \phi)|y, r^k, \phi^k]$,
 M-step:
 $r^{k+1} = \tilde{\mathbb{F}}(x^k - u^k)$
 $x^{k+1} = \mathbb{G}(r^{k+1} + u^k)$
 $u^{k+1} = u^k + (r^{k+1} - x^{k+1})$
 $\phi^{k+1} = \underset{\phi}{\text{argmin}} \left\{ \tilde{f}_1(r, \phi; r^k, \phi^k) + f_3(\phi) \right\}$
}

Figure 3. Notional JPEM concept for computing the joint MAP estimates of r and ϕ using both the Plug and Play and EM frameworks. Here, $\tilde{\mathbb{F}}$ is a modified inversion operation obtained by the substitution $\tilde{f}_1(r, \phi; r^k, \phi^k) \rightarrow f_1(r, \phi)$ in the operator \mathbb{F} . A standard GDA is used to evaluate \mathbb{G} .

4.1 Inversion Operation

In the inversion operator, $\tilde{\mathbb{F}}(x^k - u^k)$, there is no coupling between the elements of r . Therefore, we use iterative coordinate descent (ICD) to update each element, r_i , independently of the others with a closed form solution. The 1D cost function for the i^{th} element is given by

$$c(r_i; \mu_i^k, C_{i,i}^k, x^k, u^k) = \log r_i + \frac{1}{r_i} (C_{i,i}^k + |\mu_i^k|^2) + \frac{1}{2\sigma_2^2} [r_i^2 - 2r_i(x_i^k - u_i^k)]. \quad (8)$$

Here, the posterior mean, μ , and covariance matrix, C , are computed as part of the construction of $\tilde{f}_1(r, \phi; r^k, \phi^k)$, as described in App. B. Note that C^k and μ^k depend on y , r^k , and ϕ^k . To minimize Eq. (8), we differentiate with respect to r_i , set the right-hand side equal to zero, and multiply both sides r_i^2 . The result is a 3^{rd} order polynomial given by

$$0 = \alpha_1 r_i^3 + \alpha_2 r_i^2 + \alpha_3 r_i + \alpha_4, \quad (9)$$

where

$$\alpha_1 = \frac{1}{\sigma_2^2}, \quad \alpha_2 = \frac{-(x_i^k - u_i^k)}{\sigma_2^2}, \quad \alpha_3 = 1, \quad \alpha_4 = -(C_{i,i} + |\mu_i|^2). \quad (10)$$

Thus, we update each element, r_i , by simply rooting Eq. (9). The cubic equation provides a general closed-form solution for the roots of a 3^{rd} order polynomial. If multiple real-valued roots exist, we use Eq. (8) to determine the root with the lowest cost.

4.1.1 Denoising Operation

As mentioned above, the denoising operation, $\mathbb{G}(r^{k+1} + u^k)$, is mathematically equivalent to the MAP estimation process with a Gaussian noise model and prior model given by $\beta f_2(x)$. Therefore, we use a GDA to reduce the cost function associated with the minimization over x in Eq. (6). Specifically, we treat $r^{k+1} + u^k$ as a noisy image and apply a GDA to obtain x^{k+1} .

Most GDAs require that we specify the noise variance. For the P&P framework, the noise variance corresponds to $\sigma_n^2 = \beta \sigma_\lambda^2$. Different values of σ_n^2 produce outputs with different characteristics. Thus, we can view σ_n^2 as a means to change the prior model of the GDA. In order to reduce the number of parameters, we set the value of σ_λ^2 automatically during initialization and use β to adjust σ_n^2 to tune the model.

For this work, we use the Block-Matching and 3D Filtering (BM3D) GDA, which works well for coherent image reconstruction using P&P.⁹ BM3D is a patched-based denoising algorithm which stacks similar patches and collaboratively filters them using shrinkage techniques in a transform domain.¹⁶ We downloaded the BM3D MATLAB code from Ref. 17 and used the default parameter values, with two exceptions; we set the noise profile to “vn” for very noisy and the hard thresholding block size to 4×4 .

4.2 Phase-Error Update

Following 7, we use a low-resolution Gaussian Markov random field (GMRF) to model the phase-error function, denoted $\bar{\phi}$. While the GMRF model is a relatively basic image prior, it represents the simple structure of the phase-error function well.^{7,8} The distribution for $\bar{\phi}$ is given by

$$p(\bar{\phi}) = \frac{1}{z} \exp \left\{ -\frac{1}{2\sigma_{\bar{\phi}}^2} \sum_{\{i,j\} \in \mathcal{P}} b_{i,j} |\bar{\phi}_i - \bar{\phi}_j|^2 \right\}, \quad (11)$$

where z is a normalizing constant, $b_{i,j}$ is the weight between neighboring pixel pairs ($\bar{\phi}_i$ and $\bar{\phi}_j$), \mathcal{P} is the set of all pair-wise cliques falling within the same neighborhood, and $\sigma_{\bar{\phi}}^2$ controls the variation in $\bar{\phi}$.¹³

We model and estimate the GMRF on a low-resolution grid to reduce the total number of unknowns. If we subsample a two-dimensional square phase-error function, $\phi \in \mathbb{R}^M$, by a factor n_b in both dimensions, we get $\bar{\phi} \in \mathbb{R}^{M/n_b^2}$. This approach reduces the total number of unknowns from $2M$ for (r, ϕ) , to $M(1 + 1/n_b^2)$ for $(r, \bar{\phi})$. Since the forward model, A_ϕ , requires a full-resolution representation of ϕ , we use a nearest-neighbor-interpolation to relate the two quantities according to

$$\phi = \Sigma \bar{\phi}, \quad (12)$$

where Σ is an $M \times M/n_b^2$ interpolation matrix with elements in the set $[0, 1]$. Note that each element in the low resolution grid, $\bar{\phi}_i$, corresponds to a $n_b \times n_b$ group of high-resolution elements in ϕ . We use the set, $\Omega_i \in \mathbb{R}^{n_b^2}$, to specify the indices of the high-resolution elements that correspond to $\bar{\phi}_i$.

Given our low-resolution GMRF model, the phase update for $\bar{\phi}$ becomes

$$\begin{aligned} \bar{\phi}^{k+1} &= \underset{\bar{\phi}}{\operatorname{argmin}} \left\{ \tilde{f}_1(r, \Sigma \bar{\phi}; r^k, \phi^k) + f_3(\bar{\phi}) \right\}, \\ &= \underset{\bar{\phi}}{\operatorname{argmin}} \left\{ -\frac{1}{\sigma_w^2} 2\operatorname{Re} \{ y^H A_{\Sigma \bar{\phi}} \mu^k \} + \frac{1}{2\sigma_{\bar{\phi}}^2} \sum_{\{i,j\} \in \mathcal{P}} b_{i,j} |\bar{\phi}_i - \bar{\phi}_j|^2 \right\}. \end{aligned} \quad (13)$$

To minimize Eq. (13), we use ICD and sequentially update each element of $\bar{\phi}$. For the i^{th} element, the 1D cost function is

$$c(\bar{\phi}_i; \mu^k) = -|\chi_i| \cos(\angle \chi_i - \bar{\phi}_i) + \frac{1}{2\sigma_{\bar{\phi}}^2} \sum_{j \in \partial i} b_{i,j} |\bar{\phi}_i - \bar{\phi}_j|^2, \quad (14)$$

where χ_i is a complex-valued scalar and $j \in \partial i$ is an index over neighboring low-resolution phase samples. Here, χ_i is given by

$$\chi_i = \frac{2}{\sigma_w^2} \sum_{n \in \Omega_i} y_n^* [F \mu^k]_n, \quad (15)$$

where $*$ indicates the complex conjugate. Intuitively, χ_i is the difference between the data and our best estimate of the propagated field, $F\mu$, summed over all high-resolution points corresponding to $\bar{\phi}_i$. Our estimate, $\hat{\phi}$, minimizes this difference in a manner that is also consistent with $p(\bar{\phi})$. We minimize Eq. (14) using a 1D line search over $\bar{\phi}_p \in [\bar{\phi}^{**} - \pi, \bar{\phi}^{**} + \pi]$, where $\bar{\phi}^{**}$ minimizes just the prior term. By minimizing Eq. (14), we obtain an estimate of the unwrapped phase errors.

4.3 JPEM Algorithm

Figure 4 shows the steps of the basic JPEM algorithm. We first initialize r^0 using a simple back-projection of the noisy data. Then, we automatically set the parameters σ_λ , σ_n^2 , and σ_w^2 , using r^0 and y . Note that we can adjust these parameters, as needed, using the unit-less scalars, γ , β , and ρ .

To initialize ϕ^0 , we use the basic JPEM algorithm in an iterative process developed in 6,7. Fig. 5 details the steps of this process, referred to as the phase-error initialization. The initial estimate of the phase-error vector

is simply $\phi \leftarrow 0$. Then, for a set number of outer-loop iterations, N_L , we allow the basic JPEM algorithm to run for N_K iterations (i.e., N_K iterations of the EM algorithm). Each time the JPEM algorithm is called, r is re-initialized, but the previous value of ϕ is used. Finally, after the phase-error initialization is complete, we run the basic JPEM algorithm to full completion. We define algorithm completion as the point when a metric, ϵ , falls below a threshold value of ϵ_T , where

$$\epsilon = \frac{\|r^k - r^{k-1}\|}{\|r^{k-1}\|}. \quad (16)$$

JPEM{
Inputs: $y, a, \phi^0, \gamma, \beta, \rho, \sigma_{\bar{\phi}}, n_b$, (either N_K or ϵ_T)
Outputs: $\hat{r}, \hat{\phi}$
Initialize: $r^0 = |A^H y|^{\circ 2}$, $u^0 = 0$, $\sigma_\lambda \leftarrow \frac{1}{\gamma} \sqrt{s^2(r^0)}$,
 $\sigma_n^2 = \beta \sigma_\lambda^2$, $\sigma_w^2 = \frac{1}{\rho} s^2(y)$
Updates:
for $k < N_K$ **or** $\epsilon > \epsilon_T$
E-step:
 $C^k = \mathcal{D} \left(\frac{\sigma_w^2}{1 + \frac{\sigma_w^2}{\sigma_n^2}} \right)$
 $\mu^k = C^k \frac{1}{\sigma_w^2} A_{\phi^k}^H y$,
M-step:
for $i = 1 : M$
 $r_i^{k+1} \leftarrow \text{root} \{ \alpha_1 r_i^3 + \alpha_2 r_i^2 + \alpha_3 r_i + \alpha_4 \}$
end for
 $x^{k+1} = \mathbf{BM3D}(r^{k+1} + u^k, \sigma_n^2)$
 $u^{k+1} = u^k + (r^{k+1} - x^{k+1})$
for all $i \in M/n_b^2$
 $\bar{\phi}_i^k = \underset{\bar{\phi}_i}{\text{argmin}} \{ c(\bar{\phi}_i; \mu^k) \}$
end for
end for
}

Figure 4. Basic JPEM algorithm. Here, $|\cdot|^{\circ 2}$ indicates the element-wise magnitude square of a vector and $s^2(\cdot)$ computes the sample variance of a vector's elements.¹⁸ We can tune the model using the unit-less parameters γ , β , and ρ . The function $\mathbf{BM3D}(r, \sigma_n^2)$ indicates the BM3D algorithm applied to vector, r , with noise variance σ_n^2 .

Phase-Error Initialization {
Inputs: $y, a, \gamma, \beta, \rho, \sigma_{\bar{\phi}}, n_b, N_K, N_L$
Outputs: ϕ^{N_L}
 $\phi^0 = 0$
for $j = 0 : N_L - 1$ **do**
 $\phi^{j+1} = \mathbf{JPEM} \{ y, a, \phi^j, \gamma, \beta, \rho, \sigma_{\bar{\phi}}, n_b, N_K \}$
end for
}

Figure 5. Iterative process used to initialize the phase errors. The output, ϕ^{N_L} is obtained by executing the JPEM algorithm N_L times.



Figure 6. Commonly used test images, \mathcal{I} , normalized so that $\mathcal{I} \in [0, 1]$.

5. METHODS

In this section, we explain the methods used to compare the JPEM, MBIR, and IS algorithms over range of object types and contrast values. First, we describe the data simulation process. Then, we describe how each algorithm was configured. Finally, we present the metrics used to quantitatively compare performance.

5.1 Data Generation

As inputs to our simulation, we used the ten common gray-scale images shown in Fig. 6[†]. We denote these 256×256 normalized images as \mathcal{I} and generated reflectances functions with the desired contrast according to $r = \mathcal{I}^{\alpha}$. Here, $\alpha \in [0.5, 5]$ and $^{\circ}$ indicates the application of the exponent to each element. Applying the exponent, α , to each element of the normalized image, changes the variance of the pixel distribution, relative to the peak value of one. For $\alpha < 1$, elements in the set $(0, 1)$ increase in value and the distribution narrows as the values approach one. Conversely, for $\alpha > 1$, elements in the set $(0, 1)$ decrease in value and the distribution spreads out between zero and one, thereby increasing the global contrast. To quantify global image contrast, we used the standard deviation of the pixel distribution, measured in decibels (dB). We define this metric as

$$\sigma_{dB} = \sqrt{s^2(10 \log_{10} r)}, \quad (17)$$

where $s^2(\cdot)$ computes the sample variance of a vector's elements and $r \in (0, 1]$.¹⁸

To simulate the return field for a given reflectance function, r , we first generated a reflection coefficient according to $g \sim CN(0, \mathcal{D}(r))$, as shown in Fig 2 (b). Next, we used the numerical techniques from 20 to propagate g through a single phase screen located at the pupil plane. Following 20, we generated the atmospheric phase errors using an FFT-based technique and a Kolmogorov model for the refractive-index power-spectral density (PSD). Since our goal was to examine algorithm performance solely as a function of the input image characteristics, we fixed the turbulence realization across all data. Figure 7 shows the phase screen we used. The black dashed line in Fig. 7 indicates the pupil diameter, D_a . For this simulation, the source grid size was 1.5×1.5 m, the pupil-plane grid size was 0.34×0.34 m, and the propagation distance was 2 km. Additionally, we fixed the coherence length of the phase screen, r_0 , so that the $D_a/r_0 = 10$, where $D_a = 0.34$ m.

To simulate detection by a DH system using an off-axis IPRG, we padded the 256×256 propagated field to obtain an array size of 512×512 . Next, we applied a 512×512 binary circular pupil function, a , which had a circle of ones, 256 pixels wide, centered in an array of zeros. After the pupil, we first applied a thin-lens phase function which collimated the propagated light, then applied a Fast Fourier Transform to form an image in the focal plane. Next, we mixed the image with an off-axis reference beam and detected the resultant power.¹¹ Figure 2 (d) shows an example of the resulting digital hologram. The reference-beam power was set at approximately 80% of the well depth per detector element [i.e., 80% of 5×10^4 photoelectrons (*pe*)]. We also modeled Gaussian read noise with a standard deviation of 40 *pe* and digitized the output to 12 bits.

After detection, we took an FFT of the digital hologram data and isolated the signal of interest, which was a 256×256 complex image of the signal in the pupil plane. Figure 2 (e) shows the magnitude squared of the digital hologram spectrum. We have highlighted the 256×256 signal of interest, y , with a white dotted line and we also show it after it is isolated from the spectrum in Fig. 2 (f). For all data sets, we fixed the signal-to-noise ratio (SNR) at 100 (i.e., 40 dB), where we define

$$\text{SNR} = \frac{s^2(y) - s^2(w)}{s^2(w)}. \quad (18)$$

[†]Test images were downloaded from <https://github.com/csxn/DnCNN>¹⁹




Figure 7. Phase screen used for all simulations. The black dashed line indicates the pupil diameter, D_a , of our imaging system. We designed the phase errors so that $D_a/r_0 = 10$.

5.2 Algorithms Tested

We compared the JPEM algorithm to the IS algorithm from 10, modified to estimate the phase errors on a low-resolution grid, $\bar{\phi}$. This modification provides the IS algorithm the same benefits JPEM obtains from a reduced number of unknowns, thus allowing a more fair comparison. We also found this modification necessary since the point-by-point algorithm from 10 does not work well for single-shot imaging of low-contrast scenes.

The modified IS algorithm computes the low-resolution phase-error estimate according to

$$\hat{\phi} = \operatorname{argmax}_{\bar{\phi}} \left\{ - \sum_{p,q} (|F^H \mathcal{D}(\exp \{j \Sigma \bar{\phi}\})^H y|^2)^{0.5} \right\}. \quad (19)$$

Following the process described in 10, we used 20 iterations of conjugate gradient to optimize Eq. (19) and the algorithm was initialized using a 15th order Zernike polynomial estimate of the phase errors. We obtained the polynomial estimate using an iterative method to estimate only up to the 3rd order terms, then up to the 4th, and so on, continuing up to 15th order.¹⁰ Finally, we calculated the full-resolution estimate, $\hat{\phi}$, according to Eq. (12).

In addition to the modified IS algorithm, we also compared the JPEM algorithm to the MBIR algorithm from 7. These two algorithms share many of the same parameters, which we kept consistent. Specifically, we set $N_K = 60$, $N_L = 50$, $\epsilon = 1 \times 10^{-3}$, $\sigma_{\bar{\phi}} = 0.15$, $\gamma = 2$, $\rho = 1$ for the phase-error initialization, and $\rho = 2$ for the final reconstruction. For the MBIR prior model parameters, we set $T = 0.1$, $q = 2$, and $p = 1.1$. For the JPEM prior model parameter, we set $\beta = 20$. For all three algorithms, IS, MBIR, and JPEM, we subsampled $\phi \in \mathbb{R}^{256^2}$ by a factor $n_b = 4$ to obtain $\bar{\phi} \in \mathbb{R}^{64^2}$.

5.3 Quality Metrics

To measure the distortion between the reconstructed images, \hat{r} , and the simulation input, r , we used the Peak Signal-to-Noise Ratio (PSNR) function in Matlab. The PSNR is the ratio, measured in dB, of the peak-pixel value in r , to the mean square error in the reconstruction, \hat{r} . Additionally, to measure distortion between the reconstructed phase error, $\hat{\phi}$, and the actual phase error, ϕ , we used a metric referred to here as the peak Strehl ratio, S_p . We define S_p as

$$S_p = \frac{\{ |F\mathcal{D}(a)e^{j(\hat{\phi}-\phi)}|^2 \}_{\max}}{\{ |F\mathcal{D}(a)|^2 \}_{\max}}, \quad (20)$$

where $\{\cdot\}_{\max}$ indicates that we take the maximum value of the argument. The peak Strehl ratio is a measure of the peak power in the point-spread function of the imaging system, after correction.

6. RESULTS

Figures 8-10 show example reconstructions for three different images as a function of σ_{dB} . The top row of each figure shows the true reflectance, r . The contrast of the images increase from left to right. The row below the truth images shows the basic images obtained by inverting the data, according to $\hat{r}_{basic} = |A^H y|^2$, with no phase-error correction. The next three rows show the IS, MBIR, and JPEM reconstructions, respectively.

From Figs. 8-10, we see that the IS algorithm provides only marginal improvements over the basic reconstructions. For both r_{basic} and r_{IS} , the residual phase errors and high-contrast speckle variations make object characterization difficult, even for high values of σ_{dB} , as defined by Eq. (17). The MBIR reconstructions are more focused and contain less speckle variation compared to the basic and IS reconstructions. However, the MBIR outputs still contain pixel variations which create neighborhoods of either dark or light pixels. This problem is severe for low values of σ_{dB} , but tappers off as σ_{dB} increases. Finally, the JPEM algorithm produces images which are focused and appear more smooth. For low values of σ_{dB} , the JPEM reconstructions are more recognizable than those of MBIR.

Figures 11 and 12 provide a quantitative comparison of the three algorithms by showing average PSNR and peak Strehl ratio, S_p , as a function of σ_{dB} . In this case, the average is taken over the ten different images. These figures confirm the qualitative results discussed above. With respect to image quality, we see that the JPEM algorithm produces outputs with higher PSNR than either the IS or MBIR algorithms. Above $\sigma_{dB} = 18$, the JPEM and MBIR algorithms produce similar PSNR values. However, for values below that point, JPEM is more than 2 dB higher. Finally, with respect to phase-error estimation, the JPEM algorithm produces the highest average peak Strehl ratio over the entire range of σ_{dB} values.




Figure 8. Example reconstructions for *peppers* image




Figure 9. Example reconstructions for *bird* image

7. CONCLUSION

In this paper, we developed a novel algorithm, referred to as JPEM, for jointly estimating the MAP estimates of both the reflectance, r , and the phase errors, ϕ , from coherent imaging data using advanced prior models. Our algorithm is a hybrid of the joint variable estimation approach from Ref. 7 and the modified P&P framework from Ref. 9. We designed the JPEM algorithm to better reconstruct complex gray-scale reflectance functions in the presence of isoplanatic phase errors.

To test our our algorithm, we simulated data from a DH system for a wide range of input images. Specifically, we used ten different gray-scale images and varied the image contrast. We then compared the JPEM reconstructions with those from the IS and MBIR algorithms. The results show that for complex images, contrast is an important factor that affects reconstruction quality. For very low contrast values, all three algorithms produced poor results. Performance improves as the contrast increases; however, the relative rate of of this increase was different for the various algorithms. The JPEM algorithm produced the highest-quality reflectance estimates for all but the highest contrast values, and the highest-quality phase-error estimates in all cases. As a result, the JPEM images are more focused, more smooth, and more recognizable than the images produced by IS or MBIR, especially when the contrast is low.

APPENDIX A. LOG LIKELIHOOD MODEL FOR COHERENT IMAGING

In this Appendix, we specify the negative log likelihood function for a coherent system. For a reflectance function, r , we model the reflection coefficient, g , as a complex normal conditional distribution given by

$$p(g|r) \sim CN(0, \mathcal{D}(r)), \quad (21)$$

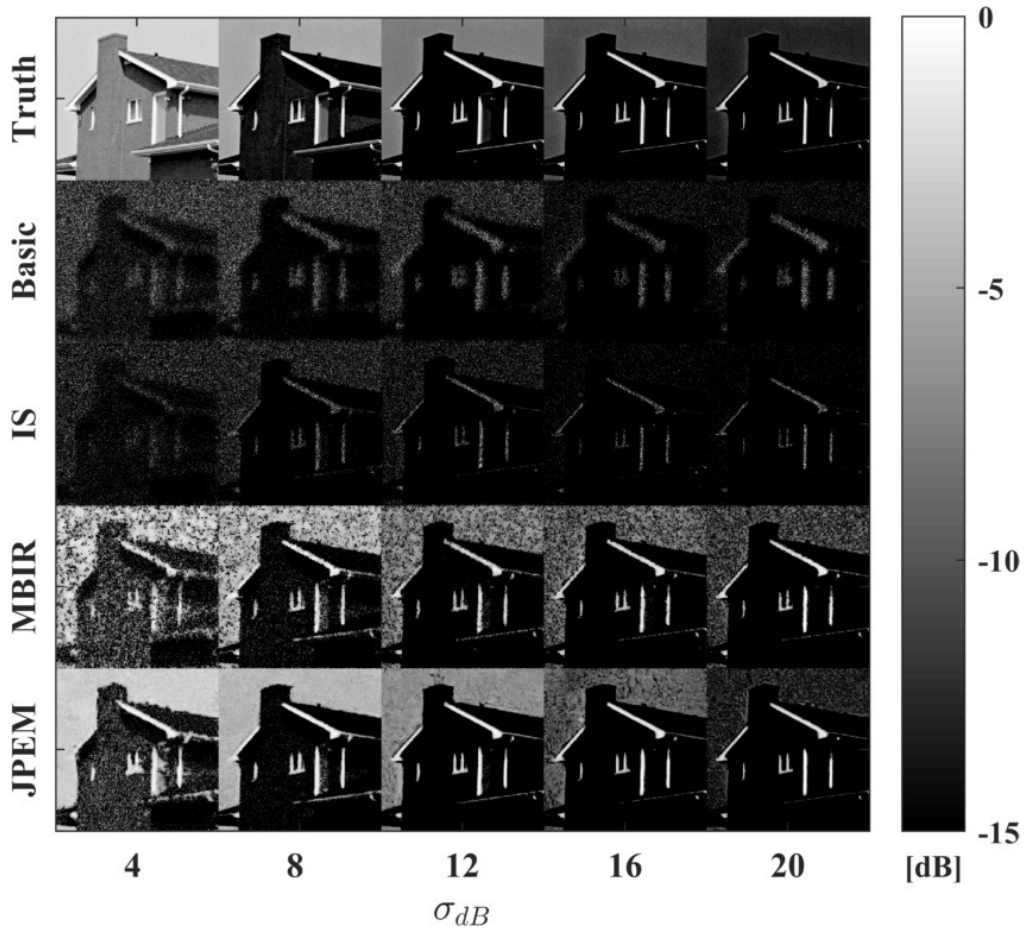


Figure 10. Example reconstructions for *house* image

where $\mathcal{D}(\cdot)$ denotes an operator that produces a diagonal matrix from its vector argument. Furthermore, we treat the detection process as shot-noise limited and represent the measurement noise, w , as additive with a zero-mean i.i.d. complex normal distribution given by

$$p(w) \sim CN(0, \sigma_w^2 I), \quad (22)$$

where σ_w^2 is the noise variance and I is the identity matrix. Given Eqs. (1) and (21-22), the resulting complex pupil image data, y , also has a complex normal conditional distribution given by⁷

$$p(y|r, \phi) \sim CN(0, A_\phi \mathcal{D}(r) A_\phi^H + \sigma_w^2 I). \quad (23)$$

The resulting negative log likelihood function for coherent systems is

$$-\log p(y|r, \phi) = \log(|A_\phi \mathcal{D}(r) A_\phi^H + \sigma_w^2 I|) + y^H (A_\phi \mathcal{D}(r) A_\phi^H + \sigma_w^2 I)^{-1} y. \quad (24)$$




Figure 11. Average PSNR vs. σ_{dB} . Here, $\langle \cdot \rangle$ indicates an average value taken over the ten different test images.

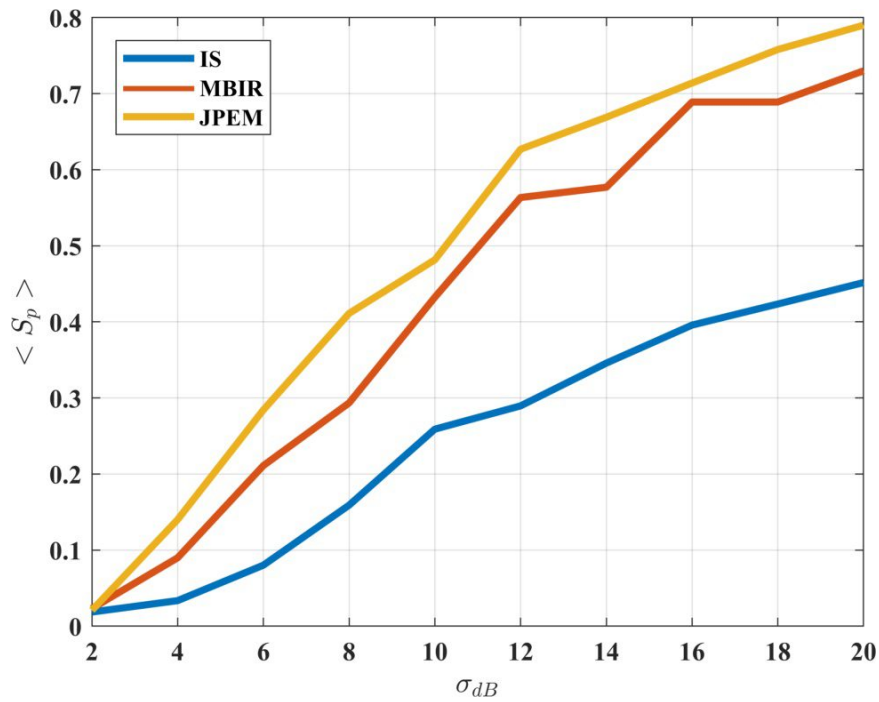


Figure 12. Average peak Strehl ratio, S_p , vs. σ_{dB} . Here, $\langle \cdot \rangle$ indicates an average value taken over the ten different test images.

APPENDIX B. CONSTRUCTION OF SURROGATE FUNCTION

In this Appendix, we show how the surrogate, \tilde{f}_1 is constructed. To begin we use Bayes' theorem and the fact that $p(y|g, r, \phi) = p(y|g, \phi)$ to write Eq. (7) as

$$\begin{aligned}\tilde{f}_1(r, \phi; r^k, \phi^k) &= -\text{E} [\log p(y, g|r, \phi)|y, r^k, \phi^k], \\ &= -\text{E} [\log p(y|g, \phi) + \log p(g|r) |y, r^k, \phi^k], \\ &= -\text{E} \left[\frac{1}{\sigma_w^2} \|y - A_\phi g\|^2 - \log |\mathcal{D}(r)| - g^H \mathcal{D}(r)^{-1} g |y, r^k, \phi^k \right] + \kappa, \\ &= -\frac{1}{\sigma_w^2} 2\text{Re} \{y^H A_\phi \mu^k\} + \log |\mathcal{D}(r)| + \sum_{i=1}^M \frac{1}{r_i} (C_{i,i}^k + |\mu_i^k|^2) + \kappa,\end{aligned}\quad (25)$$

where κ is the sum of terms constant with respect to r and ϕ , $\mu^k = \text{E} [g|y, r^k, \phi^k]$, and $C^k = \text{E} [\|g - \mu^k\|^2|y, r^k, \phi^k]$. The subscripts on μ and C indicate the vector and matrix indices, respectfully. Following 7, we find μ^k and C^k by specifying the posterior distribution of g conditioned on knowledge of y , r , and ϕ . The resulting distribution parameters are given by

$$\mu^k = C^k \frac{1}{\sigma_w^2} A_{\phi^k}^H y \quad (26)$$

and

$$\begin{aligned}C^k &= \left[\frac{1}{\sigma_w^2} A_{\phi^k}^H A_{\phi^k} + \mathcal{D}(r^k)^{-1} \right]^{-1}, \\ &\approx \mathcal{D} \left(\frac{\sigma_w^2}{1 + \frac{\sigma_w^2}{r^k}} \right).\end{aligned}\quad (27)$$

Here, we dramatically simplify the computation required to construct \tilde{f}_1 by approximating $A_\phi^H A_{\phi'} \approx I$. With the exception of $\mathcal{D}(a)$, all matrices that make up A_ϕ in Eq. (2) are unitary. Therefore, when $\mathcal{D}(a)^H \mathcal{D}(a) = \mathcal{D}(a) = I$, the approximation is exact. Intuitively, this occurs when the aperture is square, a criteria that can be artificially achieved by selecting a square subset of the pupil image, y .⁸ In practice, we have found that approximating $\mathcal{D}(a)^H \mathcal{D}(a) = I$ works well, even for circular apertures.

REFERENCES

- [1] Lucke, R. L. and Rickard, L. J., "Photon-limited synthetic aperture imaging for planet surface studies," *Applied Optics* **41**(24), 5,084–5,095 (2002).
- [2] Beck, S. M., Buck, J. R., Buell, W. F., Dickinson, R. P., Kozlowski, D. A., Marechal, N. J., and Wright, T. J., "Synthetic-aperture imaging laser radar: laboratory demonstration and signal processing," *Applied Optics* **44**(35), 7621–7629 (2005).
- [3] Buell, W., Marechal, N., Buck, J., Dickinson, R., Kozlowski, D., Wright, T., and Beck, S., "Demonstration of synthetic aperture imaging LADAR," in *[Defense and Security]*, 152–166, International Society for Optics and Photonics (2005).
- [4] Krause, B., Buck, J., Ryan, C., Hwang, D., Kondratko, P., Malm, A., Gleason, A., and Ashby, S., "Synthetic aperture LADAR flight demonstration," in *[CLEO: Science and Innovations]*, PDPB7, Optical Society of America (2011).
- [5] Poon, T.-C. and Liu, J.-P., *[Introduction to modern digital holography: with MATLAB]*, Cambridge University Press (2014).
- [6] Pellizzari, C., Trahan III, R., Zhou, H., Williams, S., Williams, S., Nemati, B., Shao, M., and Bouman, C. A., "Synthetic aperture LADAR: A model based approach," *Submitted to IEEE Transactions on Computational Imaging* (TBD).
- [7] Pellizzari, C., Spencer, M. F., and Bouman, C. A., "Phase-error estimation and image reconstruction from digital-holography data using a bayesian framework," *JOSA A* **34**(9), 1659–1669 (2017).

- [8] Pellizzari, C. J., Banet, M. T., Spencer, M. F., and Bouman, C. A., “Demonstration of single-shot digital holography using a bayesian framework,” *J. Opt. Soc. Am. A* **35**, 103–107 (Jan 2018).
- [9] Pellizzari, C., Trahan III, R., Zhou, H., Williams, S., Williams, S., Nemati, B., Shao, M., and Bouman, C. A., “Optically coherent image formation and denoising using plug and play inversion framework,” *Applied Optics* **56** (June 2017).
- [10] Thurman, S. T. and Fienup, J. R., “Phase-error correction in digital holography,” *JOSA A* **25**(4), 983–994 (2008).
- [11] Spencer, M. F., Raynor, R. A., Banet, M. T., and Marker, D. K., “Deep-turbulence wavefront sensing using digital-holographic detection in the off-axis image plane recording geometry,” *Optical Engineering* **56**(3), 031213–031213 (2017).
- [12] Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J., “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends® in Machine Learning* **3**(1), 1–122 (2011).
- [13] Bouman, C. A., “Model based image processing.” <https://engineering.purdue.edu/bouman/publications/pdf/MBIP-book.pdf> (2013).
- [14] Venkatakrisnan, S. V., Bouman, C. A., and Wohlberg, B., “Plug-and-play priors for model based reconstruction,” in [*Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE*], 945–948, IEEE (2013).
- [15] Sreehari, S., Venkatakrisnan, S., Wohlberg, B., Drummy, L. F., Simmons, J. P., and Bouman, C. A., “Plug-and-play priors for bright field electron tomography and sparse interpolation,” *arXiv preprint arXiv:1512.07331* (2015).
- [16] Dabov, K., Foi, A., Katkovnik, V., and Egiazarian, K., “Image denoising by sparse 3-D transform-domain collaborative filtering,” *IEEE Transactions on image processing* **16**(8), 2080–2095 (2007).
- [17] Dabov, K., “Image and video denoising by sparse 3D transform-domain collaborative filtering.”
- [18] Forbes, C., Evans, M., Hastings, N., and Peacock, B., [*Statistical distributions*], John Wiley & Sons (2011).
- [19] Zhang, K., Zuo, W., Chen, Y., Meng, D., and Zhang, L., “Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising,” *IEEE Transactions on Image Processing* **26**(7), 3142–3155 (2017).
- [20] Schmidt, J. D., [*Numerical Simulation of Optical Wave Propagation, With examples in Matlab*], SPIE (2010).