

Imaging through distributed-volume aberrations using single-shot digital holography

CASEY J. PELLIZZARI,^{1,*}  MARK F. SPENCER,² AND CHARLES A. BOUMAN³

¹Department of Physics, United States Air Force Academy, USAFA AFB, Colorado 80840, USA

²Air Force Research Laboratory, Directed Energy Directorate, Kirtland AFB, Albuquerque, New Mexico 87111, USA

³School of Electrical and Computer Engineering, Purdue University, West Lafayette, Indiana 47907, USA

*Corresponding author: casey.pellizzari@usafa.edu

Received 21 September 2018; revised 27 November 2018; accepted 27 November 2018; posted 28 November 2018 (Doc. ID 346294); published 2 January 2019

This paper explores the use of single-shot digital holography data and a novel algorithm, referred to as multiplane iterative reconstruction (MIR), for imaging through distributed-volume aberrations. Such aberrations result in a linear, shift-varying or “anisoplanatic” physical process, where multiple-look angles give rise to different point spread functions within the field of view of the imaging system. The MIR algorithm jointly computes the maximum *a posteriori* estimates of the anisoplanatic phase errors and the speckle-free object reflectance from the single-shot digital holography data. Using both simulations and experiments, we show that the MIR algorithm outperforms the leading multiplane image-sharpening algorithm over a wide range of anisoplanatic conditions. © 2019 Optical Society of America

<https://doi.org/10.1364/JOSAA.36.000A20>

Provided under the terms of the [OSA Open Access Publishing Agreement](#)

1. INTRODUCTION

Digital holography (DH) uses coherent illumination and spatial-heterodyne detection to sense the amplitude and phase of light scattered off an object's surface [1]. The resulting data are sensitive to phase errors caused by index-of-refraction perturbations in, for example, the atmosphere or optical systems. With this in mind, we can estimate these phase errors directly from the DH data for wavefront sensing purposes or to digitally correct aberrated images [2–6].

Image-sharpening (IS) algorithms can estimate the phase errors from DH data by maximizing an image-sharpness metric. In practice, the sharpened images result from estimates of the complex-valued reflection coefficient, g , which is the complex-valued ratio of the reflected field to the incident field. For surfaces that are rough relative to the illumination wavelength, this estimation process leads to images with high-spatial-frequency variations known as speckle. Therefore, IS algorithms require incoherent averaging of multiple speckle realizations to accurately estimate the phase errors—a process referred to here as multishot DH [2].

Recently, we developed a model-based iterative reconstruction (MBIR) algorithm for jointly computing the maximum *a posteriori* (MAP) estimates of the phase errors, ϕ , and the real-valued reflectance, r , from a single data realization—a process referred to here as single-shot DH [5]. We define the reflectance as $r = E[|g|^2]$, where $E[\cdot]$ indicates the expected value [7]. In general, r is smoother and has higher spatial correlation when compared to the reflection coefficient, g .

Reconstructing r , rather than g , allows us to leverage its higher spatial correlation to better constrain the estimation process and produce more accurate estimates of the phase errors with fewer data and less signal, compared to IS algorithms [5,6].

Both the IS algorithm in Ref. [2] and the MBIR algorithm in Ref. [5] were designed for the estimation and correction of phase errors that are isoplanatic. In general, isoplanatic phase errors (IPEs) result in a shift-invariant point spread function (PSF) in the image domain. We typically model IPEs using a single phase screen located near the entrance pupil of the imaging system [8]. This model accurately represents scenarios where the phase errors are concentrated near the observer, such as the imaging of space-based objects through atmospheric turbulence along mostly vertical propagation paths.

For other scenarios, the phase errors are often distributed along the propagation path, such as the imaging of Earth-based objects through atmospheric turbulence along mostly horizontal propagation paths. In these cases, the phase errors are anisoplanatic, meaning they result in a shift-variant PSF in the image domain. We typically model anisoplanatic phase errors (APEs) using a discrete representation consisting of multiple phase screens distributed along the propagation path [8]. To help quantify APE effects, we use the isoplanatic angle, θ_0 [9]. This atmospheric turbulence parameter, in practice, gives us a gauge for the image-domain angular separation over which the PSF is essentially unchanged. Therefore, points in the image separated by an angle larger than θ_0 have significantly different PSFs due to the APEs.

Importantly, the IS and MBIR algorithms in Refs. [2] and [5] can only correct for a single PSF across the entire image field of view (FOV). Thus, in the presence of APEs, these algorithms cannot obtain diffraction-limited performance. Estimating and correcting APEs using DH data is difficult due to increased model complexity and an increased number of unknowns. Extensions of the basic IS algorithm from [2] have been demonstrated for the correction of APEs [10–12]. In particular, Tippie and Fienup demonstrated an IS algorithm for estimating APEs modeled as multiple phase screens distributed along the propagation path [11]. However, this IS algorithm still requires multishot DH data to accurately estimate the phase errors and form a focused speckle-free image.

In this paper, we propose a novel algorithm for estimating APEs and reconstructing focused speckle-free images from single-shot DH data. Our algorithm, referred to as multiplane iterative reconstruction (MIR), uses a Bayesian framework to jointly estimate both the APEs modeled as multiple phase screens distributed along the propagation path, along with the object’s reflectance, r . Using MIR, we compute the joint MAP estimates of each phase screen along with the focused image. We compare the proposed algorithm to the multiplane IS (MIS) algorithm found in Ref. [11] over a range of anisoplanatic conditions using both simulated and experimental data. The results show that the MIR algorithm can more accurately estimate APEs and form focused, speckle-reduced images from single-shot DH data as compared to the leading MIS algorithm.

In what follows, we first formulate the overall concept for the MIR algorithm and review the prior models used in our Bayesian framework. Next, we formulate the execution of the MIR algorithm and our methods for quantifying performance. Results then follow with discussion and a conclusion.

2. MIR ALGORITHM

Figure 1 shows an example DH system using an off-axis image-plane recording geometry (IPRG) [4]. The system images an object with a reflectance function, r , and a corresponding

reflection coefficient, g , through distributed-volume phase errors, ϕ . In this example, atmospheric turbulence causes anisoplanatic effects in the DH data; however, the analysis is equally applicable to more general types of APEs, as we show in the experimental results section. For mathematical simplification, it is common to treat distributed-volume phase errors, ϕ , as a finite number of discrete layers, $[\phi_1, \phi_2, \dots, \phi_K]$, where each layer is a thin unit-amplitude phase screen that represents the phase perturbations along an incremental volume [8]. After encountering ϕ , the signal is imaged onto a detector array, where it is mixed with a reference field to form the digital hologram.

Figure 2 provides insight into the conventional image-processing steps used for a DH system with an off-axis IPRG. We first extract the signal, y , from the spectrum of the digital hologram. Due to the discrete Fourier transform (DFT) relationship between the pupil plane and the image/detector plane, y is a complex image of the signal at the pupil plane. Thus, the conventional approach to form an image is to take an inverse fast Fourier transform (FFT) of the data, y . Before moving on in the analysis, the reader should note that we provide additional details about this DH IPRG model in Appendix B.

In this work, our goal is to jointly compute the MAP estimates of the reflectance, $r \in \mathbb{R}^M$, and the phase errors, $\phi = [\phi_1, \phi_2, \dots, \phi_K]$, where $\phi_k \in \mathbb{R}^M \forall k$, from a single noisy data realization, $y \in \mathbb{C}^M$. Here, we use vectorized notation for all 2D variables. The joint MAP estimates are given by

$$\begin{aligned}
 (\hat{r}, \hat{\phi}) &= \underset{r, \phi}{\operatorname{argmin}} \{-\log p(r, \phi | y)\} \\
 &= \underset{r, \phi}{\operatorname{argmin}} \{-\log p(y | r, \phi) - \log p(r) - \log p(\phi)\}, \quad (1)
 \end{aligned}$$

where r and ϕ are assumed to be independent. Here, $p(\cdot)$ and $p(\cdot | \cdot)$ represent conventional and conditional probability distributions, respectively. For coherent imaging applications, the log-likelihood function, $\log p(y | r, \phi)$, is not easily evaluated due to its computational complexity [13]. Therefore, we require a simplified approach to solve Eq. (1). In the remaining subsections, we derive the form of $\log p(y | r, \phi)$, develop a

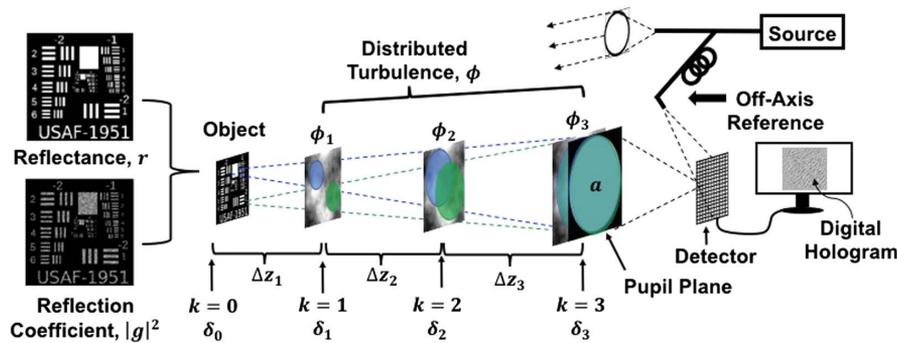


Fig. 1. Example DH system using an off-axis IPRG. A laser source is used to flood-illuminate an object that has a reflectance function, r , and corresponding reflection coefficient, g . In this example, the return signal is corrupted by atmospheric turbulence, which causes APEs, ϕ . We model ϕ as a layered atmosphere consisting of multiple phase screens, $\phi = [\phi_1, \phi_2, \dots, \phi_K]$, where $K = 3$ in this example. In our discrete model, each propagation plane, k , has an associated grid sample spacing, δ_k , and the distance between grids is given by Δz_i . The blue and green dotted cones represent ray traces from two different points on the object. The ray traces show that light from different points on the object will encounter different phase errors, resulting in anisoplanatic conditions. The returning signal passes through the pupil-aperture function, a , is collimated, and is imaged onto a focal-plane detector array, where it is mixed with a reference field to form the digital hologram.

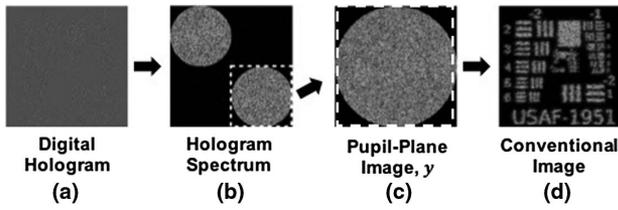


Fig. 2. Conventional image-processing steps for a DH system using an off-axis IPRG. Here, we show the magnitude squared of any complex-valued quantities. To extract the signal, we start with (a) a real-valued digital hologram, then take a 2D FFT to obtain the (b) complex-valued holographic spectrum. Note that the digital hologram, shown in (a), naturally has low contrast. Next, a section of the spectrum is isolated (c), 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 (d).

simplified framework using the expectation maximization (EM) algorithm, and define our distributions for r and ϕ .

A. Log-Likelihood Model for Coherent Imaging

For a DH system using an off-axis IPRG, we model the complex pupil image,

$$y = A_\phi g + w, \quad (2)$$

as a linear transform of the reflection coefficient, $g \in \mathbb{C}^N$, plus complex-valued measurement noise, $w \in \mathbb{C}^M$ [5]. In Ref. [5], Eq. (2) was derived for a DH system using an off-axis pupil-plane recording geometry. However, as shown in Appendix B, the model also holds for the pupil image, y , obtained from a DH system using an off-axis IPRG. The transform, $A_\phi \in \mathbb{C}^{M \times M}$, accounts for the propagation and measurement geometry and is dependent on ϕ .

For anisoplanatic conditions, the estimation framework of Eq. (1) and data model of Eq. (2) are more complex than they are for the isoplanatic case treated in Ref. [5]. First, we must account for intermediate propagations between multiple phase screens, which adds complexity to the structure of A_ϕ . With this in mind, we use the split-step beam propagation method (BPM) to model A_ϕ as a series of Fresnel propagations between the planes of interest, as described in Appendix A [8]. Second, we must model and estimate higher dimensional phase errors, where for K phase screens, we get $\phi \in \mathbb{R}^{KM}$. This makes the joint estimation of r and ϕ significantly more challenging, since the number of unknowns increases relative to the measurements.

To determine the likelihood function, $p(y|r, \phi)$, we first define distributions for g and w ; then we relate these distributions to y using Eq. (2). For a surface with reflectance, r , which is rough relative to the illumination wavelength, we model the vector g as

$$p(g|r) \sim CN(0, \mathcal{D}(r), 0), \quad (3)$$

where $\mathcal{D}(\cdot)$ denotes an operator that produces a diagonal matrix from its vector argument and $CN(\mu, C, \Gamma)$ indicates a multivariate complex normal distribution with mean, μ , covariance

matrix, C , and pseudo-covariance matrix, Γ [5]. Here, $\Gamma = 0$, since the phase of g is uniformly distributed. Next, to model the stochastic nature of coherent detection, we treat the measurement process as being shot-noise limited and use an additive, zero-mean, complex Gaussian model [14]. The distribution for w is therefore given by

$$p(w) \sim CN(0, \sigma_w^2 I, 0), \quad (4)$$

where σ_w^2 is the noise variance. Finally, using the data model from Eq. (2) and the distributions from Eqs. (3) and (4), the likelihood function for y , given r and ϕ , becomes

$$p(y|r, \phi) \sim CN(0, A_\phi \mathcal{D}(r) A_\phi^H + \sigma_w^2 I, 0), \quad (5)$$

where the superscript H indicates the Hermitian transpose.

B. Prior Models

Both the reflectance, r , and the phase errors, ϕ_k , are modeled using Markov random fields (MRFs) with the form

$$p(\Delta) = \frac{1}{z} \exp \left\{ - \sum_{\{i,j\} \in \mathcal{P}} b_{i,j} \rho(\Delta) \right\}, \quad (6)$$

where Δ is the difference between neighboring pixel pair values, z is the partition function, $b_{i,j}$ is the weight between pixel pairs, \mathcal{P} is the set of all pair-wise cliques falling within the same neighborhood, and $\rho(\cdot)$ is the potential function [15]. We obtain different models for the reflectance and phase errors by choosing different potential functions.

1. Reflectance

For the reflectance function, r , we used a Q-generalized Gaussian Markov random field (QGGMRF) potential function with the form

$$\rho_r \left(\frac{\Delta_r}{\sigma_r} \right) = \frac{|\Delta_r|^p}{p \sigma_r^p} \left(\frac{\left| \frac{\Delta_r}{T \sigma_r} \right|^{q-p}}{1 + \left| \frac{\Delta_r}{T \sigma_r} \right|^{q-p}} \right), \quad (7)$$

where T is a unitless threshold value that controls the transition of the potential function from having the exponent q to having the exponent p , and σ_r controls the variation in \hat{r} [16]. The parameters of the QGGMRF potential function affect its shape, and therefore the influence neighboring pixels have an effect on one another. In general, the QGGMRF potential function is designed to preserve edges.

2. Phase Errors

We assume that the layers of the distributed-volume phase error, ϕ , are statistically independent [8]. Therefore, we can write the distribution for ϕ as the product of the distributions for the individual screens, such that

$$p(\phi) = \prod_{k=1}^K p(\phi_k). \quad (8)$$

Estimating phase-error functions for multiple planes can significantly increase the number of unknowns that must be estimated relative to the number of measurements, M . For K screens, we must estimate KM unknowns for ϕ , plus M unknowns for r , giving a total of $M(K+1)$ unknowns. To reduce the number of unknowns, we model the phase errors,

ϕ_k , using a low-resolution Gaussian Markov random field (GMRF), denoted as $\bar{\phi}_k$ [5]. If ϕ_k is subsampled by a factor of n_b in both dimensions to obtain $\bar{\phi}_k$, each screen will have M/n_b^2 unknowns. The value of n_b can be a function of the propagation plane, k . When the same value is used for all planes, this technique reduces the total number of unknowns to $M(K/n_b^2 + 1)$.

The Gaussian potential function for $\bar{\phi}_k$ is given by

$$\rho_{\bar{\phi}_k} \left(\frac{\Delta_{\bar{\phi}_k}}{\sigma_{\bar{\phi}_k}} \right) = \frac{|\Delta_{\bar{\phi}_k}|^2}{2\sigma_{\bar{\phi}_k}^2}, \quad (9)$$

where $\Delta_{\bar{\phi}_k}$ is the difference between neighboring phase samples and $\sigma_{\bar{\phi}_k}$ controls the assumed variation in $\bar{\phi}_k$. This potential function enforces a strong influence between neighboring samples, which generates a more smoothly varying output when compared to the QGGMRF model.

While the phase errors are estimated as a low-resolution function, $\bar{\phi}$, we must up-sample the estimate to the resolution of ϕ when applying the forward-model operator, A_ϕ . Our interpolation scheme is given by

$$\phi_k = P\bar{\phi}_k, \quad (10)$$

where P is an $M \times M/n_b^2$ nearest-neighbor-interpolation matrix with elements in the set $\{0, 1\}$. More complex interpolations can be used; however, nearest neighbor allows for fast computation.

C. MIR Framework

Using Eqs. (5)–(10), we can write the MAP cost function from Eq. (1) as

$$\begin{aligned} c(r, \bar{\phi}) &= -\log p(y|r, \bar{\phi}) - \log p(r) - \log p(\bar{\phi}) \\ &= \log \left| A_{\bar{\phi}} \mathcal{D}(r) A_{\bar{\phi}}^H + \sigma_w^2 I \right| + y^H (A_{\bar{\phi}} \mathcal{D}(r) A_{\bar{\phi}}^H + \sigma_w^2 I)^{-1} y \\ &\quad + \sum_{\{i,j\} \in \mathcal{P}} b_{i,j} \rho_r \left(\frac{\Delta_r}{\sigma_r} \right) + \sum_{k=1}^K \sum_{\{i,j\} \in \mathcal{P}} \rho_{\bar{\phi}_k} \left(\frac{\Delta_{\bar{\phi}_k}}{\sigma_{\bar{\phi}_k}} \right). \end{aligned} \quad (11)$$

As noted in Ref. [5], the determinant and inverse in Eq. (11) make it difficult to evaluate; however, we may use the EM algorithm to replace Eq. (11) with a more tractable surrogate function.

The EM algorithm represents a special case of majorization minimization using a surrogate function. For any function $c(x)$, we define a surrogate function, $Q(x; x')$, to be an upper-bounding function, such that $c(x) \leq Q(x; x') + \kappa$, where x' is the current value of x that determines the functional form of Q , and κ is a constant that ensures the two functions are equal at x' . Surrogate functions have the property that minimization of $Q(x; x')$ implies minimization of $c(x)$ [15].

The EM algorithm provides a formal framework for developing a surrogate function [17]. If we choose g to be the unobserved data, our surrogate function becomes

$$\begin{aligned} Q(r, \bar{\phi}; r', \bar{\phi}') &= -E_g [\log p(y, g|r, \bar{\phi}) | y, r', \bar{\phi}'] \\ &\quad - \log p(r) - \log p(\bar{\phi}), \end{aligned} \quad (12)$$

where the conditional expectation is taken over g , given the data, y , the current estimate of the reflectance, r' , and the

current estimate of the phase errors, $\bar{\phi}'$. By introducing g as the unobserved data, we create a cost function that is more tractable than Eq. (11) because of the linear relationship between the data, y , and the reflection coefficient, g . Taking the expectation helps overcome the fact that we do not know g . Evaluation of the expectation in Eq. (12), with respect to the random vector, g , constitutes the E-step of the EM algorithm. This step establishes the specific functional form of Q . Next, we conduct the M-step to jointly minimize Q according to

$$(\hat{r}, \hat{\bar{\phi}}) = \underset{r, \bar{\phi}}{\operatorname{argmin}} \{Q(r, \bar{\phi}; r', \bar{\phi}')\}. \quad (13)$$

The EM algorithm consists of iteratively applying the E- and M-steps shown in Eqs. (12) and (13) and updating r' and $\bar{\phi}'$. For the optimization during the M-step, we use iterative coordinate descent (ICD) to sequentially minimize the entire surrogate cost function with respect to each element, r_s and $\bar{\phi}_{k,s}$, where the joint subscripts k, s indicate the s th sample of the k th phase screen [15]. Figure 3 shows the steps of the MIR algorithm. In Appendix D, we provide the update equations for each step. The update for each reflectance element, r_s , has a closed form solution and the update for each phase-error element, $\bar{\phi}_{k,s}$, requires a 1D line search. Notice that the separable form of Q derived in Appendix C ensures that sequential optimization of r and $\bar{\phi}$ is equivalent to joint optimization of these two quantities.

Since $Q(r, \bar{\phi}; r', \bar{\phi}')$ is a surrogate function for $c(r, \bar{\phi})$, the alternating steps of the EM algorithm converge in a stable manner to the local minima of the original MAP cost function given by Eq. (11) [13]. Furthermore, since both the MAP cost function and its surrogate are nonconvex, the algorithm is likely to converge to a local minimum that is not global. Since we cannot directly evaluate the MAP cost function to compare different local minima, we cannot easily select among different local minima based on the values of their MAP cost function. Thus, to ensure we find a good local minimum, in the sense that it produces a focused and useful image, we pay close attention to our initial conditions, and we employ an iterative initialization process first developed in Ref. [13]. For a set number of outer-loop iterations, N_L , we allow the basic MIR algorithm to run for N_K iterations. Each time the MIR algorithm is called, r is reinitialized, but the previous value of $\bar{\phi}$ is used. After the outer loop runs N_L times, we again run the basic

MIR{

Repeat:

E-step:

$$Q(r, \bar{\phi}; r^i, \bar{\phi}^i) = -E_g [\log p(y, g|r, \bar{\phi}) | y, r^i, \bar{\phi}^i] - \log p(r) - \log p(\bar{\phi})$$

M-step:

for all $s \in S$

$$r_s^{i+1} = \underset{r_s}{\operatorname{argmin}} \{Q(r, \bar{\phi}; r^i, \bar{\phi}^i)\}$$

for all $k \in K$

for all $s \in \bar{S}$

$$\bar{\phi}_{k,s}^{i+1} = \underset{\bar{\phi}_{k,s}}{\operatorname{argmin}} \{Q(r^i, \bar{\phi}; r^i, \bar{\phi}^i)\}$$

}

Fig. 3. MIR algorithm for the MAP estimates of r and $\bar{\phi}$. Here, S is the set of all samples in r , K is the number of partial-propagation planes, \bar{S} is the set of all samples in the low-resolution phase screens, and i is the iteration index.

MIR algorithm. For this final reconstruction run, we stop the algorithm when a metric, ϵ , falls below a threshold value of ϵ_T , where

$$\epsilon = \frac{\|r^i - r^{i-1}\|}{\|r^{i-1}\|}. \quad (14)$$

3. RESULTS

In this section, we demonstrate the utility of our proposed MIR algorithm using both simulated and experimental data. In both cases, we compare the performance of MIR to the MIS algorithm from [11] over a range of anisoplanatic conditions. Additionally, we include reconstructions from the MBIR algorithm [5], which assumes IPEs. For the reader's convenience, we provide an overview of the MIS algorithm in Appendix E.

A. Simulation

To generate synthetic data, we simulated the scenario shown in Fig. 1 in which atmospheric turbulence causes APEs. We used the split-step BPM MATLAB code from [8] to simulate the propagation of light from the object to the pupil plane. For atmospheric turbulence, we used three Kolmogorov phase screens of equal strength, distributed over the propagation path. We quantified the strength of the distributed-volume phase errors using two parameters, $D/r_{0,sw}$ and $\tilde{\theta}_0 = \theta_0/(\lambda/D)$. The first parameter, $D/r_{0,sw}$, is the ratio of the pupil diameter, D , to the spherical-wave coherence length, $r_{0,sw}$ [9]. High values of $D/r_{0,sw}$ correspond to strong phase errors, and low values correspond to weak phase errors. The second parameter, $\tilde{\theta}_0$, is the isoplanatic angle, θ_0 , normalized by the diffraction-limited viewing angle, λ/D . As previously noted, θ_0 provides a gauge for the image-domain angular separation over which the PSF is essentially unchanged [9]. Thus, high values of $\tilde{\theta}_0$ indicate neighboring points in the image have similar PSFs. As $\tilde{\theta}_0 \rightarrow 1$, points that are just barely resolved will have different PSFs, making reconstruction difficult. After propagating the fields to the pupil plane, we simulated the detection process for a DH system using an off-axis IPRG. The signal-to-noise ratio (SNR) was fixed for all data sets, and in Appendix F, we provide additional details about the simulations.

To numerically quantify algorithm performance in simulation, we used two metrics, normalized root mean square error (NRMSE) and peak Strehl ratio. NRMSE measures the distortion between the reconstructed images, \hat{r} , and the simulation input, r . It is given by

$$\text{NRMSE} = \sqrt{\frac{\|\alpha^* \hat{r} - r\|^2}{\|r\|^2}}, \quad (15)$$

where

$$\alpha^* = \underset{\alpha}{\text{argmin}}\{\|\alpha \hat{r} - r\|^2\} \quad (16)$$

is the least-squares fit for any multiplicative scalar offset between r and \hat{r} .

The peak Strehl ratio, S_p , measures the quality of our phase-error estimate, $\hat{\phi}$. It is important to note that we should not simply compare $\hat{\phi}$ to the simulated phase errors, ϕ , directly. The imaging system is underdetermined, and there are

infinitely many solutions that will produce the same effects. It is possible to obtain an estimate, where $\hat{\phi} \neq \phi$, that can still conjugate the effects of the distributed-volume phase errors and form a focused image. Therefore, to assess the quality of the reconstructed phase errors, we measure how well $\hat{\phi}$ conjugates the effects of ϕ during backpropagation. Specifically, we backpropagate the pupil function, a , with uniform phase, from $k = 3$ to the object plane at $k = 0$, through a set of phase screens, $\phi_c = \phi - \hat{\phi}$. Our estimate, $\hat{\phi}$, should ideally cancel out the effects of ϕ during backpropagation. We define the peak Strehl ratio as

$$S_p = \frac{[|A_{\phi_c}^H \mathcal{D}(a)|^2]_{\max}}{[|A_0^H \mathcal{D}(a)|^2]_{\max}}, \quad (17)$$

where A_0^H indicates the backpropagation through a vacuum and $[\cdot]_{\max}$ indicates that we take the maximum value of the argument. In Eq. (17), the numerator is the maximum intensity in the object plane that occurs when we backpropagate the function a , using $\hat{\phi}$ to conjugate the effects of ϕ at each partial-propagation plane. The denominator represents the maximum intensity that occurs in the object plane when there are no turbulence effects (i.e., backpropagation through a vacuum). Thus, S_p quantifies how close we are to the diffraction limit, ignoring the effects of tilt. Note that this Strehl-ratio metric only measures the phase-error conjugation for a single propagation angle; however, we observed that other angles have similar degrees of correction. Furthermore, we found this Strehl-ratio metric to be more useful at assessing the quality of $\hat{\phi}$ compared to image-quality metrics, such as those used in Ref. [11].

Using the methods described above, we generated 10 independent and identically distributed (i.i.d.) data realizations for 10 different values of $\tilde{\theta}_0 \in [100, 1]$, and three different values of $D/r_{0,sw} \in \{5, 10, 15\}$. Each i.i.d. data realization had unique turbulence, speckle, and noise realizations. Figures 4–6 show example reconstructions. The images displayed were

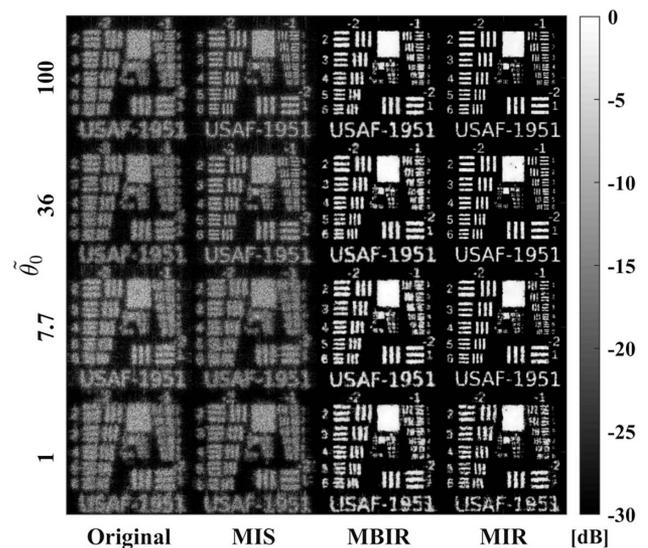


Fig. 4. Example single-shot reconstructions (simulated) for $D/r_{0,sw} = 5$. Images are displayed on a log-based decibel (dB) scale obtained according to $r_{\text{db}} = 10 \log\{r/\max(r)\}$, where $\max(\cdot)$ indicates the maximum value in the 2D array.

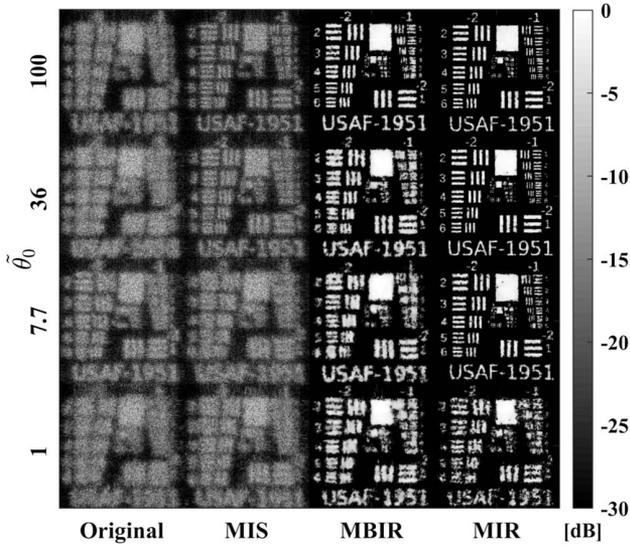


Fig. 5. Example of single-shot reconstructions (simulated) for $D/r_{0,sw} = 10$.

those with the median value of S_p from the 10 i.i.d. realizations at each value of $\tilde{\theta}_0$ and $D/r_{0,sw}$. We also show the original uncompensated image in the left column. Figures 7 and 8 show S_p and NRMSE, averaged over the 10 i.i.d. realizations, as a function of $\tilde{\theta}_0$, for each $D/r_{0,sw}$. We include the Strehl ratio for the uncompensated phase errors, computed using Eq. (17) with $\hat{\phi} = 0$, and the NRMSE for the original uncompensated image obtained according to $\hat{r} = |A_{\phi=0}^H|^2$.

The results show that the polynomial-based MIS algorithm is able to correct aberrations in the less challenging conditions when $D/r_{0,sw}$ is low and $\tilde{\theta}_0$ is high. For a fixed value of $D/r_{0,sw}$, the MIS algorithm’s performance tapers off as $\tilde{\theta}_0$ decreases. For fixed values of $\tilde{\theta}_0$, MIS performance falls off quickly as $D/r_{0,sw}$ increases. Similar to MIS, for a fixed value of $D/r_{0,sw}$, the MIR performance tapers off as $\tilde{\theta}_0$ decreases. However, the MIR

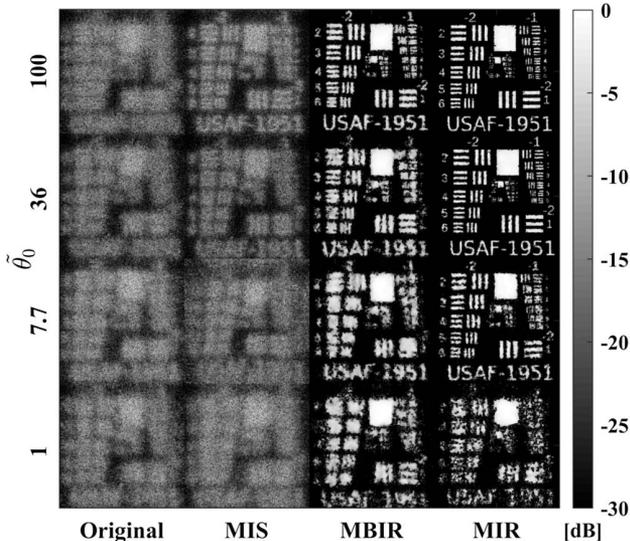


Fig. 6. Example of single-shot reconstructions (simulated) for $D/r_{0,sw} = 15$.

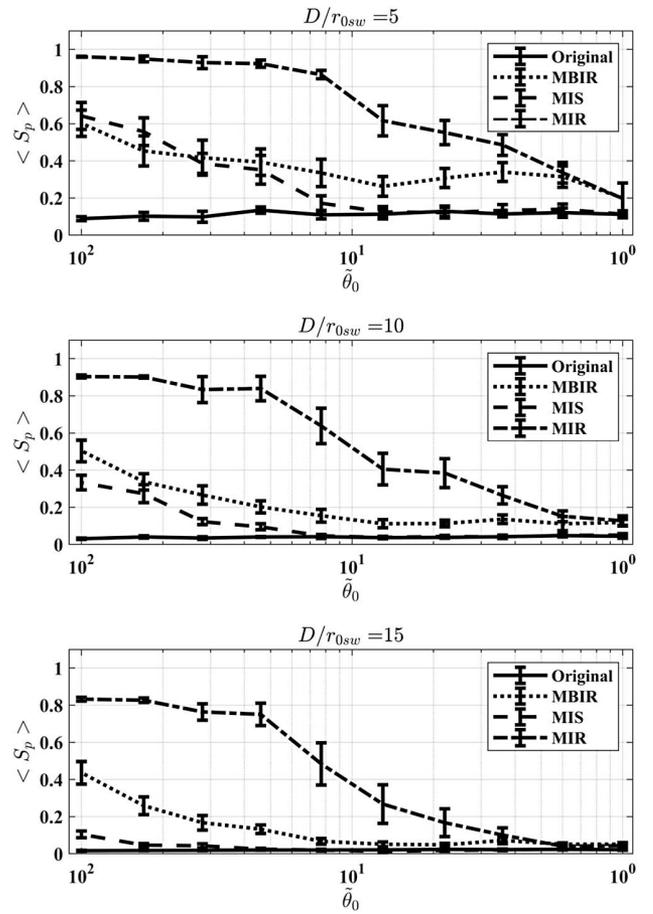


Fig. 7. Average peak Strehl ratio, S_p , as a function of $\tilde{\theta}_0$ for $D/r_{0,sw} = 5$ (top), $D/r_{0,sw} = 10$ (middle), and $D/r_{0,sw} = 15$ (bottom). The error bars show the standard deviation at each point.

algorithm is able to achieve much higher peak Strehl ratios and lower NRMSE than the MIS algorithm in all scenarios. The significantly lower NRMSE produced by the MIR algorithm can be attributed to both a better estimate of $\hat{\phi}$ and a reduction in speckle variation, as seen in Fig. 4. Unlike the MIS algorithm, MIR performance does not fall off as rapidly as $D/r_{0,sw}$ increases for fixed values of $\tilde{\theta}_0$. Lastly, we see that in most cases, the MBIR algorithm, which again assumes IPEs, performs as well or better than the MIS algorithm in terms of average Strehl ratio and outperforms the MIS algorithm in terms of average NRMSE. When visually compared to the MIR algorithm, the MBIR reconstructions look similar in the weaker turbulence cases and differ more drastically in the stronger turbulence cases. Upon close inspection of the weaker turbulence cases, we see more irregularities in the MBIR reconstructions than in the MIR reconstructions. This is due to the large difference in average Strehl ratios of the two algorithms, as shown in Fig. 7. The MBIR algorithm is simply not able to produce as accurate estimates of the APEs as the MIR algorithm. This large difference is not as apparent in the images, since some of the energy that gets spread out by the turbulence is removed during the MBIR regularization process.

To assess algorithm run time, we measured S_p as a function of time for two different cases. In Case 1, the turbulence was

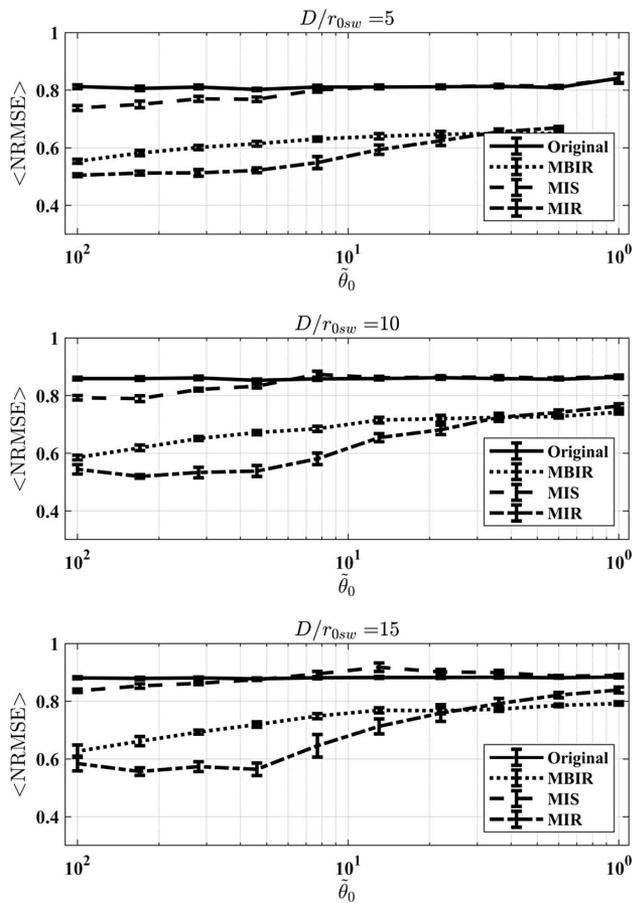


Fig. 8. Average NRMSE as a function of $\tilde{\theta}_0$ for $D/r_{0,sw} = 5$ (top), $D/r_{0,sw} = 10$ (middle), and $D/r_{0,sw} = 15$ (bottom). The error bars show the standard deviation at each point.

relatively weak, with $D/r_{0,sw} = 5$ and $\tilde{\theta}_0 = 100$. For Case 2, the turbulence was strong, with $D/r_{0,sw} = 15$ and $\tilde{\theta}_0 = 2.8$. In both cases, we used MATLAB on a PC with a 2.2 GHz Intel Xeon processor. Figure 9 plots the results for both the MIS and MIR algorithms, and Fig. 10 shows the intermediate estimates, \hat{r} , as a function of time for Case 1. The plots show that in general, the MIR algorithm achieves higher values of S_p much faster than the MIS algorithm. Additionally, we see that, for Case 1, the MIR estimate changes very little in the last

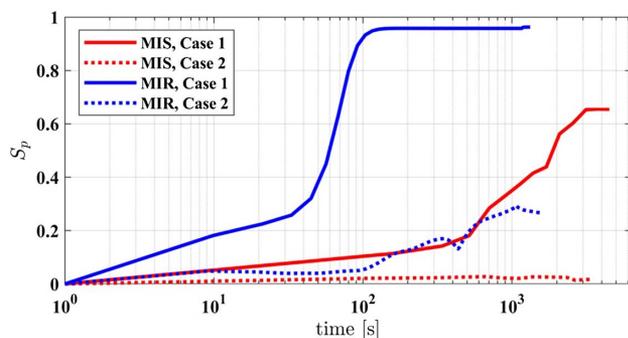


Fig. 9. Peak Strehl ratio as a function of run time for both weak (blue) and strong (red) APEs.

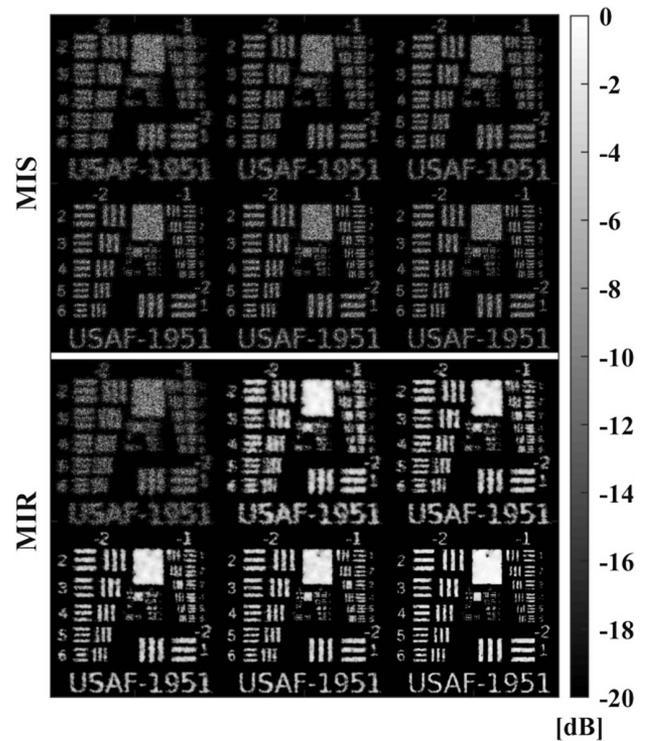


Fig. 10. Example of the intermediate samples of \hat{r} taken at various times during the algorithm run time corresponding to Case 1 in Fig. 9. The top images are the MIS estimates at $t = 0, 341, 706, 2086, 3134,$ and 4501 s. The bottom images are the MIR estimates at $t = 0, 9.9, 21.4, 33.2, 163,$ and 1324 s (ordered from left to right first, then top to bottom).

90% of the total run time; thus, we can shorten the run time and still get similar results. However, as noted in Appendix F, for simplicity we fixed the number MIR and MBIR outer-loop iterations at $N_L = 50$ and the inner-loop iterations at $N_K = 20$ for all cases. During the last outer loop, we let the algorithms run until the stopping criteria were met. On average, MIS ran 60 iterations in 67 min, MBIR ran 2565 iterations in 11 min, and MIR ran 2674 iterations in 28 min.

B. Experiment

To demonstrate the utility of the MIR algorithm with real data, we also conducted a laboratory experiment, as outlined in Appendix F. A laser with wavelength $\lambda = 1064$ nm was used to illuminate an Air Force resolution chart 2.5 m away. To simulate APEs over the propagation path, we introduced multiple phase screens made from plastic CD jewel cases. For the experiment, we started with three screens located at $z = [2.2, 2.3, 2.4]$ m, then incrementally added additional screens at distances 2.1 m, 2.0 m, and 1.9 m. However, to simplify reconstruction, we only estimated four screens located at $z = [1.9, 2.1, 2.3, 2.5]$ m for all data. Figure 11 shows the images reconstructed from the experimental data for three, four, five, and six plastic phase screens. The left column shows the original images reconstructed with no phase-error mitigation. Due to the nonuniformity of the plastic phase screens used, we see anisotropic (not to be confused with anisoplanatic)

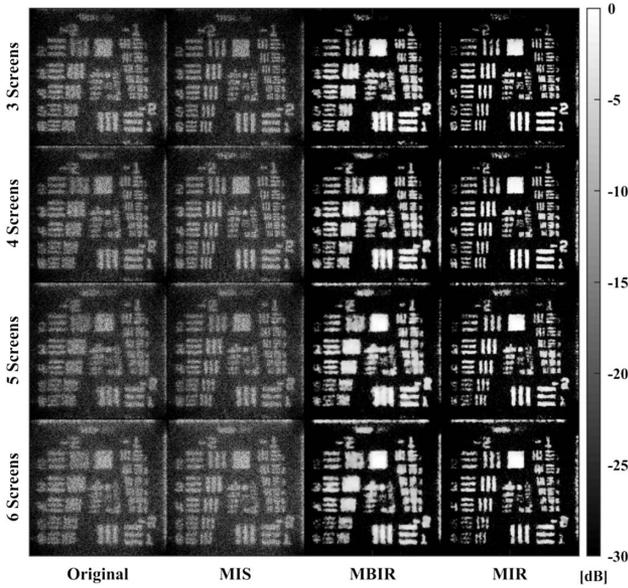


Fig. 11. Experimental results for three, four, five, and six plastic phase screens. The original reconstruction with no phase-error compensation is shown in the left column, the MIS reconstructions are in the second column, and the MIR reconstructions are in the third column. Additionally, we show reconstructions from the isoplanatic MBIR algorithm.

phase-error effects in the images. Here, the multiplane phase errors encountered in the experiment had stronger effects in the x axis than in the y axis. In Fig. 11, we also see that the phase errors were stronger in some areas of the image than others, making for a more challenging reconstruction problem.

Since the underlying statistics of the plastic phase screens used in the experiment were unknown, we cannot easily quantify performance in terms of atmospheric turbulence parameters such as $D/r_{0,sw}$ and $\tilde{\theta}_0$. However, we can compare against the simulated results to get a general idea of the experimental conditions. Doing so, we see that the overall turbulence strength most closely resembles our simulated data when $D/r_{0,sw} = 5$. Furthermore, including the reconstructions from the MBIR algorithm provides a visual representation of the degree of anisoplanatic effects, since it assumes IPEs.

Figure 11 shows that the MBIR algorithm is not able to fully correct the multiplane phase errors encountered in the experiment. Alternatively, the MIS algorithm is able to improve upon the original image in all cases, although the sharpness is reduced as the number of screens increases. Similar to the simulations, the MIR algorithm is able to produce sharper and higher-contrast images with less speckle. Note that while the experimental phase errors were clearly anisoplanatic and anisotropic, the conditions were not as challenging as those produced in the simulation.

4. CONCLUSION

In summary, we have developed a novel algorithm for jointly estimating APEs and an object's reflectance from single-shot DH data using a Bayesian framework. Our MIR algorithm models the APEs using multiple phase screens distributed along

the propagation path and computes the MAP estimates of each phase screen along with the focused reflectance image.

To test our algorithm, we designed a simulation and experiment to produce data representative of a DH system using an off-axis IPRG. We generated data for a range of anisoplanatic conditions and compared the MIR algorithm to the polynomial-based MIS algorithm from [11]. To assess the results, we presented both qualitative and quantitative comparisons.

In all cases, we observed that the polynomial-based MIS algorithm works moderately well for the less severe APEs (e.g., when the value of $\tilde{\theta}_0$ is high and $D/r_{0,sw}$ is low). However, the performance of the MIS algorithm falls off quickly as the APEs worsen (i.e., with decreases in $\tilde{\theta}_0$ and increases of $D/r_{0,sw}$). Additionally, we observed that the MIR algorithm produces high-quality images and phase-error estimates over a wide range of anisoplanatic conditions. We therefore believe the MIR algorithm provides a robust alternative to the MIS algorithm for estimating APEs and for producing focused, speckle-reduced images from single-shot DH data.

APPENDIX A: PROPAGATION MODEL

In Eq. (2), the linear transform, A_ϕ , represents the propagation of light from the object to the observer through a number of intermediate planes containing phase screens. Using the split-step BPM found in Ref. [8], we model A_ϕ as a series of partial propagations between consecutive planes. This technique uses the angular spectrum method (ASM) to model each partial propagation followed by the application of a phase screen. In this section we first review the ASM, then we use the split-step BPM to provide an explicit definition for A_ϕ in Eq. (A9).

1. ASM

To model the vacuum propagation of a monochromatic field, $\tilde{u}_0(x_0, y_0)$, at some arbitrary distance, Δz , we use the ASM [18]. Here, we use tildes to indicate continuous functions. The ASM computes the convolutional form of Fresnel diffraction integral by multiplying the input field in the spatial-frequency domain with the free-space transfer function. Note that the free-space transfer function is given by

$$\tilde{H}(f_{x_0}, f_{y_0}) = e^{j\frac{\pi}{\lambda}\Delta z} e^{-j\pi\lambda\Delta z(f_{x_0}^2 + f_{y_0}^2)}, \quad (\text{A1})$$

where λ is the wavelength and f_{x_0} and f_{y_0} are the spatial frequency coordinates [18]. Thus, at a distance Δz away, the diffraction pattern from $\tilde{u}_0(x_0, y_0)$ is given by

$$\tilde{u}_1(x_1, y_1) = \text{CSFT}^{-1}[\tilde{H}(f_{x_0}, f_{y_0})\text{CSFT}[\tilde{u}_0(x_0, y_0)]], \quad (\text{A2})$$

where $\text{CSFT}[\cdot]$ and $\text{CSFT}^{-1}[\cdot]$ indicate the continuous-space Fourier transform (CSFT) operation and its inverse, respectively.

In practice, we leverage the discrete version of Eq. (A2) found in Ref. [8]. Here, we write that discrete model, using vector notation, as

$$u_1 = F^{-1}HFu_0, \quad (\text{A3})$$

where $u_0 \in \mathbb{C}^M$ is the vectorized input field, $u_1 \in \mathbb{C}^M$ is the vectorized output field, and $F \in \mathbb{C}^{M \times M}$ is a 2D DFT matrix. The diagonal matrix, H , is the matrix form of the free-space transfer function and is given by

$$H = \mathcal{D} \left(\text{vec} \left[e^{-j\pi \frac{\Delta z}{N^2 \delta_0^2} (p^2 + q^2)} \right] \right). \quad (\text{A4})$$

Here, $\text{vec}[\cdot]$ is an operator that forms a 1D vector from a 2D argument, δ_0 is the sample spacing, N is the number of samples in each dimension of the 2D transfer function, and $p, q \in [-N/2, N/2 - 1]$ are indices in the spatial-frequency domain. For simplicity, we have dropped any complex scalars and assumed a square sampling grid where $N \times N = M$.

We can only use Eq. (A3) when the input- and output-field sample spacing and grid lengths are the same (i.e., $\delta_1 = \delta_0$). Since this is restrictive, we use a modified version of the ASM found in Ref. [8] to control the ratio $\beta_1 = \delta_1/\delta_0$, and thus adjust the output grid size, $\delta_1 N$. The modified propagation model is given by

$$u_1 = \Lambda_1 F^{-1} H_{0,1} F \Lambda_0 u_0, \quad (\text{A5})$$

where

$$\begin{aligned} \Lambda_0 &= \mathcal{D} \left(\text{vec} \left[e^{j \frac{\pi \beta_1}{\lambda \Delta z} \delta_0^2 (m^2 + n^2)} \right] \right), \\ \Lambda_1 &= \mathcal{D} \left(\text{vec} \left[e^{j \frac{\pi \beta_1 - 1}{\lambda \beta_1 \Delta z} \delta_1^2 (m^2 + n^2)} \right] \right), \end{aligned} \quad (\text{A6})$$

and

$$H_{0,1} = \mathcal{D} \left(\text{vec} \left[e^{-j \frac{\pi \Delta z}{N^2 \delta_1 \delta_0} (p^2 + q^2)} \right] \right). \quad (\text{A7})$$

Here, Λ_0 and Λ_1 are diagonal matrices that apply quadratic phase factors, $\beta_1 = \delta_1/\delta_0$ is the ratio of the sample spacing in each grid, and $m, n \in [-N/2, N/2 - 1]$ are indices in the spatial domain. We add the subscript 0,1 to indicate that $H_{0,1}$ is the free-space-transfer function between planes $k = 0$ and $k = 1$.

2. Split-Step BPM Model

The split-step BPM model combines multiple partial propagations between phase screens using the modified ASM approach described by Eq. (A5). It accounts for differences between the object-plane sample spacing, δ_0 , and the pupil-plane sample spacing, δ_K . If we convert the split-step BPM operators found in Ref. [8] to vector notation, we can write the field passing through the pupil in the K th propagation plane as

$$u_K = A_\phi u_0, \quad (\text{A8})$$

where

$$\begin{aligned} A_\phi &= \mathcal{D}(a) \Lambda_K \left[\prod_{k=1}^K \mathcal{D}(e^{j\phi_k}) F^{-1} H_{k-1,k} F \right] \Lambda_0, \\ \Lambda_0 &= \mathcal{D} \left(\text{vec} \left[e^{j \frac{\pi \beta_1}{\lambda \Delta z} \delta_0^2 (m^2 + n^2)} \right] \right), \\ \Lambda_K &= \mathcal{D} \left(\text{vec} \left[e^{j \frac{\pi \beta_K - 1}{\lambda \beta_K \Delta z} \delta_K^2 (m^2 + n^2)} \right] \right), \end{aligned} \quad (\text{A9})$$

and

$$H_{k-1,k} = \mathcal{D} \left(\text{vec} \left[e^{-j \frac{\pi \Delta z_k}{N^2 \delta_{k-1} \delta_k} (p^2 + q^2)} \right] \right). \quad (\text{A10})$$

Here, δ_k is the sample spacing in the k th plane, N is the number of samples in each dimension of the 2D functions, $m, n \in [-N/2, N/2 - 1]$ are indices in the spatial domain,

and $p, q \in [-N/2, N/2 - 1]$ are indices in the spatial-frequency domain. Additionally, in Eq. (A9), ϕ_k is the phase screen in the k th plane, Δz_k is the partial-propagation distance between plane $(k - 1)$ and plane k , and $\beta_k = \delta_k/\delta_{k-1}$. We add the subscript $k, k - 1$ to indicate that $H_{k-1,k}$ is the free-space-transfer function from planes $k - 1$ to k . As previously noted, $\mathcal{D}(a)$ is a diagonalization of the pupil-aperture vector, a , and F is a 2D DFT. Note that we have dropped any complex scalars and assumed a square sampling grid for simplicity (i.e., $N \times N = M$).

In Eq. (A9), we see that in order to model the propagation using A_ϕ , we must have knowledge of the geometry. Specifically, we must know the total distance, Δz , and the sample spacings, δ_0 and δ_K . In practice, these values can be determined from knowledge of the object range, sensor FOV, and pupil diameter. The partial-propagation distances can be an arbitrary set, depending on how many screens are to be modeled and their desired locations. Following [8], the values of δ_1 through δ_{K-1} are set automatically from the values of δ_0 and δ_K .

APPENDIX B: MEASUREMENT MODEL

In this section, we derive a discrete input–output relationship for a DH system using an off-axis IPRG. We start with the complex signal image, $\tilde{u}_S(x, y)$, a continuous quantity located at the focal-plane array (FPA) of the DH sensor. For the off-axis IPRG, we perform spatial-heterodyne detection by mixing the signal image with a reference field from a local oscillator (LO) that is split off from the master-oscillator laser. For a point source LO, located at point $(x_p = r_x, y_p = r_y)$ in the pupil plane, the reference field at the FPA is given by

$$\tilde{u}_R(x, y) = \tilde{R}_0 e^{j2\pi(f_{r_x}x + f_{r_y}y)}, \quad (\text{B1})$$

where \tilde{R}_0 is the amplitude, assumed to be spatially uniform. For brevity, we ignore the quadratic phase factor, $e^{j\frac{\pi}{\lambda f}(x^2 + y^2)}$, applied to the reference, $u_R(x, y)$, by the imaging lens. Since this phase is applied to both the reference and the signal, it cancels out of Eq. (B2). In Eq. (B1), $f_{r_x} = r_x/(\lambda f)$ and $f_{r_y} = r_y/(\lambda f)$ are the spatial frequencies associated with the modulation from the off-axis reference, and f is the focal length of the imaging lens. The resulting hologram intensity is given by

$$\begin{aligned} \tilde{i}(x, y) &= |\tilde{u}_R(x, y) + \tilde{u}_S(x, y)|^2 \\ &= |\tilde{u}_R(x, y)|^2 + |\tilde{u}_S(x, y)|^2 + \tilde{u}_R(x, y)\tilde{u}_S^*(x, y) \\ &\quad + \tilde{u}_R^*(x, y)\tilde{u}_S(x, y), \end{aligned} \quad (\text{B2})$$

where $*$ indicates the complex conjugate. Substituting Eq. (B1) into Eq. (B2) gives

$$\begin{aligned} \tilde{i}(x, y) &= \tilde{R}_0^2 + |\tilde{u}_S(x, y)|^2 + \tilde{R}_0 \tilde{u}_S^*(x, y) e^{j2\pi(f_{r_x}x + f_{r_y}y)} \\ &\quad + \tilde{R}_0 \tilde{u}_S(x, y) e^{-j2\pi(f_{r_x}x + f_{r_y}y)}. \end{aligned} \quad (\text{B3})$$

We model the discrete and noisy digital hologram, measured from the intensity in Eq. (B3), as

$$i(m, n) = \tilde{i}(x, y) \Big|_{x = m\delta_x, y = n\delta_y} + w_R(m, n), \quad (\text{B4})$$

where δ_x and δ_y are the sample spacing on the FPA and m and n are the sample indices in x and y dimensions. Here, we have

simplified the model by ignoring the blurring effects of the finite-sized pixels. Furthermore, we assume that the DH system is shot-noise limited with noise $w_R(m, n)$, driven by the power of the reference field [14]. Thus, we model $w_R(m, n)$ as additive, zero-mean, white Gaussian noise. Substituting Eq. (B3) into Eq. (B4) gives

$$i(m, n) = R_0^2 + |u_S(m, n)|^2 + R_0 u_S^*(m, n) e^{j(\omega_m m + \omega_n n)} + R_0 u_S(m, n) e^{-j(\omega_m m + \omega_n n)} + w_R(m, n), \quad (\text{B5})$$

where R_0 is the digitized reference amplitude, and $\omega_m = 2\pi f_{x,r} \delta_x$ and $\omega_n = 2\pi f_{y,r} \delta_y$ are the modulation frequencies associated with the digitized reference field.

Following the example shown in Fig. 2, we extract the signal by first taking a DFT of the digital hologram, $i(m, n)$. The resulting hologram spectrum is given by

$$I(p, q) = \text{DFT}[i(m, n)] \\ = R_0^2 \delta_{[p,q]} + U_S(p, q) \star U_S(p, q) + R_0 U_S^*(p + p_r, q + q_r) + R_0 U_S(p - p_r, q - q_r) + W_b(p, q). \quad (\text{B6})$$

Here, $\text{DFT}[\cdot]$ is the 2D DFT operator, $\delta_{[p,q]}$ is a discrete 2D delta function, \star denotes a discrete convolution, p_r and q_r are the spatial-frequency shifts corresponding to the modulation from the off-axis reference, $W_b(p, q) = \text{DFT}[w_b(m, n)]$, and $U_S(p, q)$ is the DFT of the complex-signal image, $u_S(m, n)$. Due to the Fourier-transform relationship between the image and pupil planes, $U_S(p, q)$ is a discrete representation of the signal field in the pupil plane [4].

From Eq. (B6), we see that the modulation from the off-axis reference shifts the spectrum of the two cross terms away from the spectrum of the first two terms, which are centered at D.C. in $I(p, q)$. Thus, we can spatially extract a grid around the term of interest (fourth term) and divide by R_0 to get our signal data, such that

$$y(p, q) = U_S(p, q) + w(p, q), \quad (\text{B7})$$

where $w(p, q)$ is the extracted and scaled noise. Finally, we can represent Eq. (B7) using vector notation as

$$y = u + w, \quad (\text{B8})$$

where y , u , and w are in the set \mathbb{C}^M and are obtained by vectorizing $y(p, q)$, $U_S(p, q)$, and $w(p, q)$, respectively. To obtain our final measurement model, given by Eq. (2), we substitute the discrete version of the propagated field, obtained using Eq. (A8), into Eq. (B8) for U .

APPENDIX C: SURROGATE FUNCTION

In this section, we derive the EM surrogate for the MAP cost function. We start by writing Eq. (12) as

$$Q(r, \bar{\phi}; r', \bar{\phi}') = -E_g[\log p(y, g | r, \bar{\phi}) | y, r', \bar{\phi}'] \\ - \log p(r) - \log p(\bar{\phi}), \\ = -E_g[\log p(y | g, \bar{\phi}) + \log p(g | r) | y, r', \bar{\phi}'] \\ - \log p(r) - \log p(\bar{\phi}), \quad (\text{C1})$$

where we have used Bayes' theorem inside the expectation and the fact that $p(y | g, r, \bar{\phi}) = p(y | g, \bar{\phi})$. Next, we substitute in the likelihood and prior models from Eqs. (5)–(10). This gives

$$Q(r, \bar{\phi}; r', \bar{\phi}') = E_g \left[\frac{1}{\sigma_w^2} \|y - A_{\bar{\phi}} g\|^2 | y, r', \bar{\phi}' \right] + \log |\mathcal{D}(r)| \\ + \sum_{i=1}^M \frac{1}{r_i} E_g[\log |g_i|^2 | y, r', \bar{\phi}'] + \sum_{\{i,j\} \in \mathcal{P}} b_{i,j} \rho_r \left(\frac{\Delta_r}{\sigma_r} \right) \\ + \prod_{k=1}^K \left[\sum_{\{i,j\} \in \mathcal{P}} b_{i,j} \rho_{\bar{\phi}_k} \left(\frac{\Delta_{\bar{\phi}_k}}{\sigma_{\bar{\phi}_k}} \right) \right] + \kappa, \quad (\text{C2})$$

where κ is a constant with respect to r and $\bar{\phi}$.

In Ref. [5], we showed that the conditional posterior distribution of g is complex Gaussian with mean

$$\mu = C \frac{1}{\sigma_w^2} A_{\bar{\phi}}^H y \quad (\text{C3})$$

and covariance

$$C = \left[\frac{1}{\sigma_w^2} A_{\bar{\phi}}^H A_{\bar{\phi}} + \mathcal{D}(r')^{-1} \right]^{-1}. \quad (\text{C4})$$

Using Eqs. (C3) and (C4), we evaluate the expectations in Eq. (C2) to get the final form of our EM surrogate function. This gives us

$$Q(r, \bar{\phi}; r', \bar{\phi}') = -\frac{1}{\sigma_w^2} 2 \text{Re}\{y^H A_{\bar{\phi}} \mu\} + \log |\mathcal{D}(r)| \\ + \sum_{i=1}^N \frac{1}{r_i} (C_{i,i} + |\mu_i|^2) + \sum_{\{i,j\} \in \mathcal{P}} b_{i,j} \rho_r \left(\frac{\Delta_r}{\sigma_r} \right) \\ + \prod_{k=1}^K \left[\sum_{\{i,j\} \in \mathcal{P}} b_{i,j} \rho_{\bar{\phi}_k} \left(\frac{\Delta_{\bar{\phi}_k}}{\sigma_{\bar{\phi}_k}} \right) \right], \quad (\text{C5})$$

where μ_i is the i th element of the posterior mean and $C_{i,i}$ is the i th diagonal element of the posterior covariance.

APPENDIX D: MIR ALGORITHM

In this section, we present details for executing each step of the MIR algorithm. The discussion describes how we evaluate the surrogate function in the E-step and how we optimize it in the M-step. We also present initialization and stopping procedures.

1. E-Step: Surrogate Function Formation

The E-step forms Eq. (C5) so that it can be minimized during the M-step. Specifically, this consists of computing the posterior mean and covariance according to Eqs. (C3) and (C4). We simplify this process by approximating the posterior covariance as

$$C \approx \mathcal{D} \left(\frac{\sigma_w^2}{1 + \frac{\sigma_w^2}{\sigma_r^2}} \right), \quad (\text{D1})$$

which requires that $A_{\bar{\phi}}^H A_{\bar{\phi}} \approx I$. The matrix $A_{\bar{\phi}}$ can be decomposed into several matrices that are all unitary, with the exception of the $\mathcal{D}(a)$, which applies the pupil-aperture function. Therefore, when $\mathcal{D}(a) = I$ and/or $\mathcal{D}(a)^H \mathcal{D}(a) = I$, the approximation of Eq. (D1) is exact. Typically, this is not the case, since a is circular, but we approximate it as such to dramatically simplify computations. In practice, we have found this to work well. Alternatively, a square subset of the

DH data can be used for which $\mathcal{D}(a) = I$ and Eq. (D1) is exact [6].

2. M-Step: Surrogate Function Optimization

During the M-step, we minimize the surrogate function according to Eq. (13). Specifically, we employ iterative coordinate descent (ICD) to sequentially minimize Eq. (C5) with respect to each element, r_s and $\tilde{\phi}_{k,s}$. Here, the joint subscripts k, s indicate the s th sample of the k^{th} phase screen [15].

To update each reflectance element, r_s , we ignore all terms in Eq. (C5) that do not contain r_s . The resulting 1D cost function can be minimized with a line search over \mathbb{R}^+ ; however, a closed-form solution can be obtained using the symmetric bound method of majorization described in Ref. [15]. Using this technique, we replace the QGGMRF potential function from Eq. (7) with a quadratic surrogate function. The resulting cost function is given by

$$q_s(r_s; r', \tilde{\phi}') = \log r_s + \frac{C_{s,s} + |\mu_s|^2}{r_s} + \sum_{j \in \partial s} \tilde{b}_{s,j} (r_s - r_j)^2, \quad (\text{D2})$$

where $\tilde{b}_{s,j}$ are the modified weights between neighboring pixel pairs given by [15]

$$\tilde{b}_{s,j} = b_{s,j} \frac{|r_s - r_j|^{p-2} \left| \frac{r_s - r_j}{T\sigma_r} \right|^{q-p} \left(\frac{q}{p} + \left| \frac{r_s - r_j}{T\sigma_r} \right|^{q-p} \right)}{2\sigma_r^p \left(1 + \left| \frac{r_s - r_j}{T\sigma_r} \right|^{q-p} \right)^2}. \quad (\text{D3})$$

To obtain a closed-form solution for \hat{r}_s , we differentiate the right-hand side of Eq. (D2) with respect to r_s , set it equal to zero, and multiply both sides by r_s^2 . The solution, \hat{r}_s , can then be found by rooting the polynomial $\alpha_1 r_s^3 + \alpha_2 r_s^2 + \alpha_3 r_s + \alpha_4$, where

$$\begin{aligned} \alpha_1 &= 2 \sum_{j \in \partial s} \tilde{b}_{s,j}, & \alpha_2 &= -2 \sum_{j \in \partial s} \tilde{b}_{s,j} r_j, \\ \alpha_3 &= 1, & \alpha_4 &= -(C_{s,s} + |\mu_s|^2). \end{aligned} \quad (\text{D4})$$

Thus, we update each sample using the closed-form solution for the roots of a third-order polynomial. If multiple real roots exist, we chose the one with the lowest cost according to Eq. (D2).

While the reflectance can be updated using a closed-form solution, the phase-error update requires a line search. To compactly express the cost function for the update of the phase-error element, $\tilde{\phi}_{k,s}$, it is helpful to write the forward-model operator as

$$A_{\tilde{\phi}} = \Psi_k \mathcal{D}(e^{-jP\tilde{\phi}_k}) \Sigma_k, \quad (\text{D5})$$

where Ψ_k represents all the matrices on the left-hand side of $\mathcal{D}(e^{-jP\tilde{\phi}_k})$ in Eq. (A9), and Σ_k represents all the matrices on the right-hand side of $\mathcal{D}(e^{-jP\tilde{\phi}_k})$. Using this notation, we can write the cost function for the k th phase screen as

$$\begin{aligned} q(\tilde{\phi}_k; r', \tilde{\phi}') &= -\frac{2}{\sigma_w^2} \text{Re}\{(\Psi_k^H y) \mathcal{D}(e^{-jP\tilde{\phi}_k}) (\Sigma_k \mu)\} \\ &+ \frac{1}{2\sigma_{\tilde{\phi}_k}^2} \sum_{\{i,j\} \in \mathcal{P}} b_{i,j} |\Delta_{\tilde{\phi}_k}|^2. \end{aligned} \quad (\text{D6})$$

If we consider just the terms in Eq. (D6), which depend on the element $\tilde{\phi}_{k,s}$, we obtain the 1D cost function for that element given by

$$\begin{aligned} q(\tilde{\phi}_{k,s}; r', \tilde{\phi}') &= -|\chi| \cos(\angle\chi - \tilde{\phi}_{k,s}) \\ &+ \frac{1}{2\sigma_{\tilde{\phi}_k}^2} \sum_{j \in \partial s} b_{i,j} |\tilde{\phi}_{k,s} - \tilde{\phi}_{k,j}|^2, \end{aligned} \quad (\text{D7})$$

where

$$\chi = \frac{2}{\sigma_w^2} \sum_{i=1}^{n_b^2} [\Psi_k^H y]_{B(i)}^* [\Sigma_k \mu]_{B(i)}. \quad (\text{D8})$$

In Eq. (D7), $j \in \partial s$ is an index over neighboring low-resolution phase samples, B is the n_b^2 set of vector indices of the high-resolution phase, $\tilde{\phi}_k$, corresponding to the low-resolution element $\tilde{\phi}_{k,s}$, and $\angle\chi$ indicates the phase of the complex-valued scalar, χ . Intuitively, χ is the difference between the fields on either side of the screen $\mathcal{D}(e^{-jP\tilde{\phi}_k})$, summed over all high-resolution points that correspond to the single low-resolution point being updated.

The cost function of Eq. (D7) is the superposition of a cosine term, which has an infinite number of minima with equal cost, and a quadratic term. The global minima for the surrogate will be near the minimum of the quadratic term, bracketed on either side by the period of the cosine. Thus, we minimize Eq. (D7) using a 1D line search over $\tilde{\phi}_{k,s} \in [\tilde{\phi}_{k,s}^{**} - \pi, \tilde{\phi}_{k,s}^{**} + \pi]$, where $\tilde{\phi}_{k,s}^{**}$ minimizes the quadratic term in Eq. (D7). After repeating this for all $s \in \tilde{S}$, where \tilde{S} is the set of samples in the lower-resolution phase-error grid, we obtain the full-resolution estimate of the unwrapped phase screen, $\hat{\phi}_k$, according to Eq. (10).

3. Initialization and Stopping Criteria

Figure 12 summarizes the steps of the basic MIR algorithm. The parameters σ_r and σ_w are set automatically, although the regularization of r can be adjusted using a unitless parameter, γ . Note that while we initialize r using a simple back-projection of the noisy data, we use an iterative process to initialize ϕ , based on the heuristic developed in Refs. [5,13].

Figure 13 details the steps of this iterative initialization process, referred to as the MIR initialization. The initial estimate of the phase-error vector is simply $\phi^0 = \mathbf{0}$. Then, for a set number of outer-loop iterations, N_L , we allow the basic MIR algorithm to run for N_K iterations. Each time the MIR algorithm is called, r is reinitialized, but the previous value of ϕ is used. After the outer loop runs N_L times, we again run the basic MIR algorithm. For this final reconstruction run, we stop the algorithm when $\epsilon < \epsilon_T$, in accordance with Eq. (14). It is important to note that while the MIR algorithm has many parameters, we set most of them automatically. In the methods section, we provide settings that we have found to robustly provide good results over a wide range of anisoplanatic conditions.

APPENDIX E: MULTIPLANE IS ALGORITHM

To investigate the performance gain of the MIR algorithm over conventional techniques, we considered two different MIS algorithms that differ in their representation of the phase errors, ϕ . A 15th-order polynomial is used for each phase screen in Ref. [11], while a point-by-point representation is used in Ref. [12]. Unfortunately, neither of these algorithms was designed for single-shot data, and only the polynomial-based

```

MIR{
Inputs:  $y, a, z, \delta, \lambda, \phi^0, \gamma$ , (either  $N_K$  or  $\epsilon_T$ )
Outputs:  $\hat{r}, \hat{\phi}$ 
Initialize:  $r^0 = |A_{\phi^0}^H y|^2$ ,  $\sigma_r = \frac{1}{\gamma} \sqrt{s^2(r^0)}$ ,  $\sigma_w^2 = s^2(y)$ 
while  $i < N_K$  or  $\epsilon > \epsilon_T$  do
     $C^i = \mathcal{D} \left( \frac{\sigma_w^2}{1 + \frac{\sigma_w^2}{r^{i-1}}} \right)$ 
     $\mu^i = C^i \frac{1}{\sigma_w^2} A_{\phi^{i-1}}^H y$ ,
    for all  $s \in S$  do
         $r_s^i = \text{root} \{ \alpha_1 r_s^3 + \alpha_2 r_s^2 + \alpha_3 r_s + \alpha_4 \}$ 
    end for
    for all  $k \in K$  do
        for all  $s \in \bar{S}$  do
             $\bar{\phi}_{k,s}^i = \underset{\bar{\phi}_{k,s}}{\text{argmin}} \{ q(\bar{\phi}_{k,s}; r^{i-1}, \bar{\phi}^{i-1}) \}$ 
        end for
         $\phi_k^i \leftarrow P \bar{\phi}_k^i$ 
    end for
end while
}

```

Fig. 12. MIR algorithm for the MAP estimates of r and ϕ . Here, $|\cdot|^2$ indicates the element-wise magnitude square of a vector, and $s^2(\cdot)$ computes the sample variance of a vector's elements [19]. Additionally, $z = [\Delta z_1, \Delta z_2, \dots, \Delta z_K]$ is a vector of the partial-propagation distances, $\delta = [\delta_0, \delta_K]$ is a vector of the object and pupil-plane sample spacing, γ is a unitless parameter introduced to tune the amount of regularization in r , S is the set of all samples in r , K is the number of partial-propagation planes, \bar{S} is the set of all samples in the low-resolution phase screens, and i is the iteration index.

algorithm produced usable results. Even with the bootstrapping initialization approach, we found that the point-by-point MIS algorithm produces poor results by over-sharpening single-shot images to a few points. Conversely, the polynomial-based MIS algorithm has fewer unknowns and is naturally constrained by the 15th-order representation. We found it to produce more useful results for single-shot data. Therefore, we chose the polynomial-based algorithm in Ref. [11] as our reference for comparisons.

The polynomial-based MIS algorithm computes the phase estimate according to

$$\hat{\phi} = [Z \oplus Z \oplus \dots \oplus Z] \hat{c}, \quad (\text{E1})$$

```

MIR Initialization {
Inputs:  $y, a, z, \gamma, N_K$ 
Outputs:  $\hat{\phi}$ 
 $\phi^0 = \mathbf{0}$ 
for  $i = 1 : N_L$  do
     $\phi^i = \text{MIR} \{ y, a, z, \delta, \phi^{i-1}, \gamma, N_K \}$ 
end for
}

```

Fig. 13. Iterative process used to initialize ϕ .

where $Z \in \mathbb{R}^{M \times N_z}$ is a matrix with columns corresponding to the first N_z Zernike polynomials and \oplus indicates a matrix direct sum. The number of Z matrices direct-summed is equal to the number of phase-error planes modeled. In Eq. (E1), the value $\hat{c} \in \mathbb{R}^{3N_z}$ is an estimate of the polynomial coefficients for all three phase screens found according to

$$\hat{c} = \underset{c}{\text{argmax}} \{ \| (|A_c^H y|^2)^{\circ \beta} \|_1 - \alpha \| \mathcal{D}(W) |FA_c^H y|^2 \|_1 \}. \quad (\text{E2})$$

Here, $\|\cdot\|_1$ indicates the L1 norm, $\beta = 1.01$ is an image-sharpness parameter, α is a scalar that balances the two optimization terms, A_c indicates a dependence of A on c , and $W \in \mathbb{R}^M$ is a vector of weights that applies a high-pass filter to the magnitude squared of the back-projected image spectrum, $|FA_c^H y|^2$. The second term in Eq. (E2) is a regularization term that prevents object demagnification [11,12,20].

Following the algorithm in Ref. [11], we use conjugate gradient (CG) to optimize Eq. (E2) in an iterative manner by first estimating only the coefficients corresponding to a third-order polynomial for all the screens, then fourth, and so on, continuing up to the 15th order. We generated W by first creating an inverse rectangular function with width $0.8N$ (i.e., $(1 - \text{rect}[p/0.8N])(1 - \text{rect}[q/0.8N])$, where p and q are the spectral indices in the set $[-N/2, N/2 - 1]$; then we applied a raised-cosine transition from the inner section of zeros to the outer section of ones. Through trial and error, we found $\alpha = 1 \times 10^{-6}$ to be optimal.

APPENDIX F: METHODS

In this section, we provide details about the methods used to produce data. We begin with a full description of the wave-optics simulation used to produce synthetic data and then describe our laboratory setup used to produce experimental data.

1. Simulation Details

Figure 1 of the paper shows the 256×256 reflectance function, r , used as the simulation input. We generated a reflection coefficient according to $g \sim CN(0, \mathcal{D}(r))$. Next, we propagated g a distance, $\Delta z = 80$ m, through three phase screens, located at $z = [26.4, 52.8, 80]$ m [8]. We generated the three phase screens using a Kolmogorov model for the refractive-index power spectral density (PSD) [8].

The strength of each individual screen was parameterized using the plane-wave coherence length, $r_{0,pw}$, also known as the Fried parameter [9]. Each of the three screens had equal values of $r_{0,pw}$. We quantified the overall strength of the distributed-volume phase errors with the parameter $D/r_{0,sw}$. Note that the spherical-wave coherence length is more appropriate than the plane-wave coherence length for distributed-volume phase errors. The value of $D/r_{0,sw}$ specifies the degrees of freedom in the phase-error function at the pupil plane and allows us to quantify the cumulative effects of the distributed-volume phase errors. High values of $D/r_{0,sw}$ correspond to strong phase errors and low values correspond to weak phase errors. For our experiment, we controlled $D/r_{0,sw}$ by varying the value of $r_{0,pw} \in [3 \times 10^{-4}, 10]$ m for the three screens.

We quantified the anisoplanatic conditions with the parameter $\bar{\theta}_0 = \theta_0/(\lambda/D)$. For our simulations, we controlled $\bar{\theta}_0$ by

adjusting (λ/D) . Specifically, we adjusted the grid spacing in the object plane, $\delta_0 \in [1 \times 10^{-4}, 1.2 \times 10^{-2}]$ m, and set the corresponding pupil-plane sample grid spacing according to $\delta_3 = \lambda \Delta z \delta_0 / 256$ m. The intermediate grid spacings, δ_1 and δ_2 , were set according to the split-step BPM MATLAB code from [8].

To simulate detection by a DH system using an off-axis IPRG, we padded the propagated field to obtain an array size of 512×512 . Next, we applied a 512×512 binary circular pupil-aperture function, a , which had a circle of ones, 256 pixels wide, centered in an array of zeros. After, we first applied a thin-lens phase function that collimated the propagated light, then applied an FFT to form an image in the focal plane. Next, we mixed the image with a reference field given by Eq. (B1) and detected the resultant power. Figure 2(a) 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(b) shows the magnitude squared of the digital hologram spectrum. The 256×256 signal of interest, y , is highlighted by a white dotted line and is also shown in Fig. 2(c) after it is removed from the spectrum. The SNR was set to 100 for all data sets, where

$$\text{SNR} = \frac{s^2(A_{\phi g})}{s^2(w)}, \quad (\text{F1})$$

and $s^2(\cdot)$ is the sample-variance operator.

For simulated reconstructions using the MIR algorithm, we allowed the outer initialization loop to run $N_L = 50$ times, with $N_K = 20$ EM iterations each time. Once the iterative initialization process was complete, we set a stopping criterion of $\epsilon_T = 1 \times 10^{-4}$ and let the EM algorithm run to completion. During initialization, we used $q = 2$, $p = 2$, $T = 1$, $\gamma = 2$, and $b = G(0.8)$, where $G(\sigma)$ indicates a 3×3 Gaussian kernel with standard deviation, σ . For the final MIR reconstruction, we set $p = 1.1$, $T = 0.1$, and $b = G(0.1)$. Additionally, for the phase-error prior model parameters, we used $b = G(0.1)$, $\sigma_{\hat{\phi}} = [0.05, 0.10, 0.15]$, and set $n_b = 4$.

2. Laboratory Experiment Details

Figures 14 and 15 describe our experimental setup. For our master oscillator (MO) laser, we used a LightWave nonplanar ring oscillator with 700 mW of continuous-wave power and approximately 1 km of coherence length, with a wavelength of 1064 nm. Past a Faraday isolator and two alignment mirrors, we used a half-wave plate and polarized beam-splitting cube to create two optical legs. In the first optical leg, we coupled the beam from the MO laser into a single-mode, polarization-maintaining fiber to create an off-axis LO. This off-axis LO created a quasi-uniform reference beam with the appropriate tilt for digital-holographic detection in the off-axis IPRG. The mean strength of this reference beam was set to approximately 75% of the camera's pixel-well depth. In the

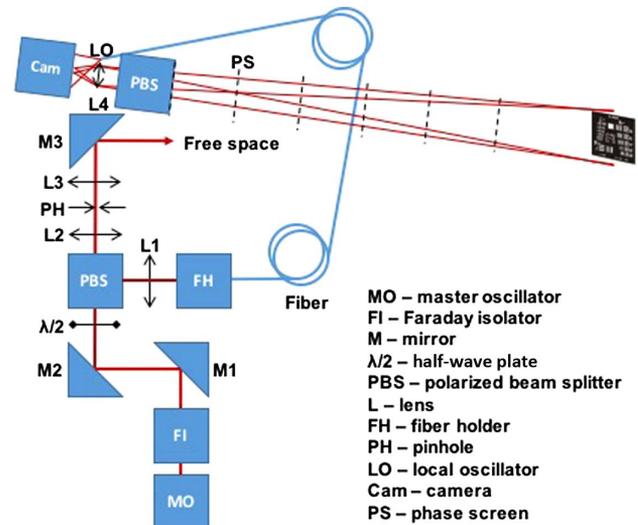


Fig. 14. Diagram of our experimental setup. The output from a 1064 nm laser was split into two paths. One path was sent through a beam expander to illuminate the object. The other path created an ideal reference that interfered with the scattered signal using digital-holographic detection in the off-axis IPRG. Phase screens were placed in front of the imaging lens to simulate the effects of anisoplanatic aberrations.

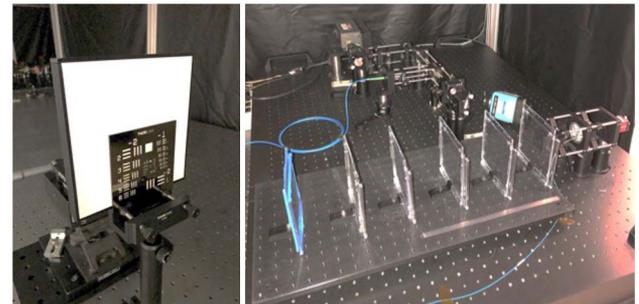


Fig. 15. Images of the object (left) and the transmitter and receiver optics (right) used in our experiment.

second optical leg, we focused the beam from the MO laser through a pinhole and then used an adjustable collimating lens to flood-illuminate the object after reflection from a steering mirror.

After setting up the imaging system with an object distance of approximately 2.5 m and an image distance of approximately 10 cm, we simulated the effects of anisoplanatic aberrations by placing six different phase screens in front of the imaging lens with a spacing of 10 cm in between. Each experimental phase screen contained two crisscrossed plastic plates (from CD jewel cases) placed back-to-back. For digital-holographic detection, we used a single lens to image the scattered light from the object onto a Guppy PRO f-125 FireWire camera. We then saved the recorded digital hologram in a 512×512 , 16-bit TIFF file. Note that our recorded digital hologram had an SNR of approximately 10 according to Eq. (F1). Also note that we positioned the object, a chrome-on-glass Air Force resolution

chart backed by Labsphere Spectralon, at an angle to minimize specular reflections and ensured that it was uniformly illuminated.

For all experimental reconstructions, MIS, MIR, and MBIR, we chose to digitally estimate four phase screens located at $z = [1.9, 2.1, 2.3, 2.5]$ m. Note that we observed better results when we included a phase screen at the aperture plane located at $z = 2.5$ m. We allowed the outer initialization loop to run $N_L = 100$ times, with $N_K = 50$ EM iterations each time. Once the iterative initialization process was complete, we set a stopping criterion of $\epsilon_T = 1 \times 10^{-4}$ and let the EM algorithm run to completion. During initialization, we used $q = 2$, $p = 2$, $T = 1$, $\gamma = 2$, and $b = G(0.8)$. For the final MIR reconstructions, we set $p = 1.1$, $T = 0.05$, $\gamma = 2.5$, and $b = G(0.1)$. Additionally, for the phase-error prior model parameters, we used $\sigma_{\hat{\phi}} = [0.15, 0.30, 0.45, 0.60]$, and set $n_b = [10, 10, 5, 5]$. For the MBIR reconstructions, we set the parameters equal to the MIR parameters, except that only the $z = 2.5$ m phase screen was reconstructed.

Acknowledgment. The authors would like to thank B. T. Plimmer and D. E. Thornton for their assistance in conducting the laboratory experiments and in reviewing this paper.

REFERENCES

1. T.-C. Poon and J.-P. Liu, *Introduction to Modern Digital Holography: With MATLAB* (Cambridge University, 2014).
2. S. T. Thurman and J. R. Fienup, "Phase-error correction in digital holography," *J. Opt. Soc. Am. A* **25**, 983–994 (2008).
3. J. C. Marron, R. L. Kendrick, N. Seldomridge, T. D. Grow, and T. A. Höft, "Atmospheric turbulence correction using digital holographic detection: experimental results," *Opt. Express* **17**, 11638–11651 (2009).
4. M. F. Spencer, R. A. Raynor, M. T. Banet, and D. K. Marker, "Deep-turbulence wavefront sensing using digital-holographic detection in the off-axis image plane recording geometry," *Opt. Eng.* **56**, 031213 (2017).
5. C. Pellizzari, M. F. Spencer, and C. A. Bouman, "Phase-error estimation and image reconstruction from digital-holography data using a Bayesian framework," *J. Opt. Soc. Am. A* **34**, 1659–1669 (2017).
6. C. Pellizzari, M. T. Banet, M. F. Spencer, and C. A. Bouman, "Demonstration of single-shot digital holography using a Bayesian framework," *J. Opt. Soc. Am. A* **35**, 103–107 (2018).
7. J. W. Goodman, *Speckle Phenomena in Optics: Theory and Applications* (Roberts and Company, 2006).
8. J. D. Schmidt, *Numerical Simulation of Optical Wave Propagation with Examples in MATLAB* (SPIE, 2010).
9. L. Andrews and R. Phillips, *Laser Beam Propagation through Random Media*, 2nd ed. (SPIE, 2005).
10. S. T. Thurman and J. R. Fienup, "Correction of anisoplanatic phase errors in digital holography," *J. Opt. Soc. Am. A* **25**, 995–999 (2008).
11. A. E. Tippie and J. R. Fienup, "Phase-error correction for multiple planes using a sharpness metric," *Opt. Lett.* **34**, 701–703 (2009).
12. A. E. Tippie and J. R. Fienup, "Multiple-plane anisoplanatic phase correction in a laboratory digital holography experiment," *Opt. Lett.* **35**, 3291–3293 (2010).
13. C. J. Pellizzari, R. Trahan, H. Zhou, S. Williams, S. E. Williams, B. Nemat, M. Shao, and C. A. Bouman, "Synthetic aperture lidar: a model-based approach," *IEEE Trans. Comput. Imaging* **3**, 901–916 (2017).
14. V. V. Protopopov, *Laser Heterodyning*, Vol. 149 of Springer Series in Optical Sciences (Springer, 2009).
15. C. A. Bouman, *Model Based Image Processing*, 2013, <https://engineering.purdue.edu/~bouman/publications/pdf/MBIP-book.pdf>.
16. J. B. Thibault, K. Sauer, C. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multi-slice helical CT," *Med. Phys.* **34**, 4526–4544 (2007).
17. A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. R. Stat. Soc. Ser. B* **39**, 1–38 (1977).
18. J. W. Goodman, *Introduction to Fourier Optics*, 3rd ed. (Roberts and Company, 2005).
19. C. Forbes, M. Evans, N. Hastings, and B. Peacock, *Statistical Distributions* (Wiley, 2011).
20. A. E. Tippie, "Aberration correction in digital holography," Ph.D. dissertation (University of Rochester, 2012).