

Multi-Agent Consensus Equilibrium for Range Compressed Holographic Surface Reconstruction

Tony G. Allen (a,b), David J. Rabb (a), Gregory T. Buzzard (b), Charles A. Bouman (b)

(a) Air Force Research Laboratory, Sensors Directorate, LADAR Technologies Branch
3109 Hobson Way, WPAFB, OH 45433

(b) Purdue University

610 Purdue Mall, West Lafayette, IN 47907

Lead Author e-mail address: tony.allen.5@us.af.mil

Abstract: Recently developed image reconstruction algorithms for range-compressed holography (RCH) have proven effective at correcting distributed-volume atmospheric phase errors. However, these methods are limited by speckle and measurement noise. In this work, we reformulate the RCH inversion problem in terms of the target's underlying speckle-free reflectance rather than the complex-valued field. We also present an algorithm for surface reconstruction from phase-corrected RCH data under the Multi-Agent Consensus Equilibrium framework, which combines advanced surface models for prior information with a physics-based model for data fidelity. Together, this reformulation and the use of advanced priors reduces the effects of speckle and leverages the spatial correlation in reflectance to suppress measurement noise. Finally, we present a point-cloud-based metric of image quality and demonstrate that our algorithm outperforms traditional RCH inversion methods both qualitatively and in terms of this metric, particularly in the low photoelectron count regime.

Keywords: Digital Holography, Image Processing, Deep Learning

1. Introduction

Digital holography (DH) is a powerful imaging technique for recovering the complex electromagnetic field reflected from a distant object. In brief, detection in a DH system involves measuring the modulation of a strong reference field by a potentially weak return field. This modulation allows for the detection to be shot-noise limited, even in scenarios with a low-power return signal. In practice, DH systems are sensitive to phase errors, whether caused by atmospheric turbulence or defects in the optical system, which need to be corrected in order to form a focused image. Moreover, many applications require DH systems to operate with limited laser power, resulting in images with significant noise and speckle. There has been great success in mitigating these effects through computational imaging algorithms. Of particular success have been a variety of Image Sharpening algorithms [1–5] which correct atmospheric phase errors by maximizing a sharpness metric, and Plug and Play (PnP) algorithms [6–10] which, on top of correcting phase errors, mitigate measurement noise and speckle variation through an extension of a Bayesian formulation that allows for the use of advanced prior models such as neural network denoisers.

As technologies have advanced, extending the methods of DH to three-dimensional (3D) imaging has been the focus of much recent work. One example of this is Range Compressed Holography (RCH), in which frequency diversity is introduced during the hologram recording [11,12]. In RCH, a series of DH images are recorded, each corresponding to a unique laser frequency. The images are then range compressed by Fourier transform over the frequency spectrum resulting in a 3D image. Relative to 2D DH, this technique faces additional challenges in terms of phase aberrations [13–16]. Recent work in RCH image formation include [14], which addresses differential phase aberrations across temporal frequencies by a phase gradient algorithm, and [15,16] which extend the Image Sharpening approach to this context. However, these methods do not address the problem of speckle and noise mitigation, which could allow RCH systems to operate with less laser power.

In this paper, we reformulate the RCH inversion problem in terms of the target's underlying speckle-free reflectance rather than the complex-valued field. We also present an algorithm for surface

reconstruction from phase-corrected RCH data under the Multi-Agent Consensus Equilibrium (MACE) [17] framework, which combines advanced surface models for prior information with a physics-based model for data fidelity. We assume no atmospheric phase errors are present and instead focus on surface detection in the presence of speckle and measurement noise. Together, this reformulation and the use of advanced priors reduces the effects of speckle and leverages the spatial correlation in reflectance to suppress measurement noise. Finally, we present a point-cloud-based metric of image quality and demonstrate that our algorithm outperforms traditional RCH inversion methods both qualitatively and in terms of this metric, particularly in the low photoelectron count regime.

2. Estimation Framework

For an opaque object illuminated by a coherent source of wavelength λ , DH allows one to image the field in the object plane,

$$U(x, y; \lambda) = f(x, y) \exp\left\{2\pi j \frac{2d(x, y)}{\lambda}\right\}, \quad (1)$$

where (x, y) are the coordinates of the object plane, $f(x, y)$ is the complex reflectance of the object, $d(x, y)$ is the distance from the object to the object plane, and the extra factor of 2 accounts for the round-trip distance. For simplicity, ignoring effects such as magnification, we assume the field in the object plane to be the ideal field recovered in the image plane. Using a single wavelength, the range information of the object is only available modulo 2π on the order of the wavelength.

By a process called range compression, one can further disambiguate the range information by introducing wavelength (frequency) diversity. This can be achieved in several ways, such as recording multiple holograms at individually tuned frequencies [15], or by using a linear frequency chirp [13].

The principal relation of RCH expresses the ideal image of the object's 3D complex reflectance, $g(x, y, z)$, as the inverse Fourier transform $U(x, y; \lambda)$ as a function of $\xi = 2/\lambda$,

$$g(x, y, z) = \int_{-\infty}^{\infty} U(x, y; \xi) \exp\{2\pi j \xi x\} d\xi = f(x, y) \delta(z - d(x, y)). \quad (2)$$

This formulation assumes the object is opaque, allowing just one range return for each (x, y) cross-range location, but it can be generalized to account for non-opaque objects as well. Under appropriate diffraction-limited, narrowband, and shot-noise-limited assumptions [19], the modulated, noisy, complex image $Y(x, y, z)$ at the image plane can be modeled as

$$Y(x, y, z) = (g * h)(x, y, z) + \eta(x, y, z), \quad (3)$$

where $*$ denotes convolution, $h(x, y, z)$ represents the point spread function (PSF) that accounts for the effects of propagation, and w is the measurement noise.

In practice, the imaged field is more accurately modeled as a discrete sampling of $Y(x, y, z)$. We can write a discrete representation of the model in Eq. (3) as $y = Ag + \eta$, where $A \in \mathbb{C}^{N \times N}$ is the matrix representation of convolution with h , $g \in \mathbb{C}^N$ is the vectorized complex reflectance, and $\eta \in \mathbb{C}^N$ is measurement noise distributed as $CN(0, \sigma_\eta^2 I, 0)$ [20]. Together, this gives the conditional distribution of our measure data $y \in \mathbb{C}^N$ given the complex reflectance g as $CN(Ag, \sigma_\eta^2 I, 0)$. Standard inversion methods produce a reconstruction of g , which leads to a speckled intensity image, $|g|^{\circ 2}$, where the modulus-squared is taken elementwise.

Instead of estimating g , we estimate its expected intensity, or speckle-free reflectance, $r = E[|g|^{\circ 2}]$. As in [8], we model speckle $g|r$ as a circularly symmetric complex Gaussian distribution. Combining this with $y|g$, we get the conditional distribution of the data given the speckle free reflectance as $y|r \sim CN\left(0, (AD(r)A^H + \sigma_\eta^2 I)^{-1}, 0\right)$, where $D(r)$ is a diagonal matrix with elements given by r .

In practice, a common method to reduce speckle is to exploit multiple recordings of the object that are captured under various realizations (looks) of the noise process. This is often done by slight adjustments to the recording set-up that provide independent noise samples while keeping the signal highly correlated. This can be effectively modeled as independent measurements of the data $y_\ell|r$ that have the same distribution as $y|r$ for each look, ℓ .

Our method of estimating r is based on the maximum a posteriori (MAP) estimate, which can be seen as a regularization of the maximum likelihood estimate. Assuming L looks, this can be written as

$$\hat{r}_{MAP} = \underset{r \in \mathbb{R}^N}{\operatorname{argmin}} \left\{ -\sum_{\ell=1}^L \log p(y_\ell | r) - \sum_{k=1}^K \beta_k h_k(r) \right\}, \quad (4)$$

where $\beta_k > 0$ controls the amount of regularization applied by h_k , and together the K regularizers constitute a prior model. The following section describes an algorithm for solving a generalization of Eq. (4) that allows for the use of state-of-the-art regularization techniques.

3. Algorithm

Under the MACE framework, the optimization problem of Eq. (4) is equivalent to a set of consensus equilibrium (CE) conditions. As in [17], we define the operators $F_\ell(w; \sigma^2)$ and $H_k(w; \sigma^2)$ to be proximal maps of $-\log p(y_\ell | r)$ and $-\beta_k h_k(r)$, respectively. Further, we define the stacked vector $\mathbf{w} = [w_1, \dots, w_{L+K}] \in \mathbb{R}^{N \times (L+K)}$ as the concatenation of the $L + K$ state vectors, along with the operators $F : \mathbb{R}^{N \times (L+K)} \rightarrow \mathbb{R}^{N \times (L+K)}$ and $G : \mathbb{R}^{N \times (L+K)} \rightarrow \mathbb{R}^{N \times (L+K)}$ as

$$F(\mathbf{w}) = [F_1(w_1), \dots, F_L(w_L), H_1(w_{L+1}), \dots, H_K(w_{L+K})], \quad (5)$$

$$G(\mathbf{w}) = [\bar{\mathbf{w}}, \dots, \bar{\mathbf{w}}], \quad (6)$$

where $\bar{\mathbf{w}} = \sum_{i=1}^{L+K} \mu_i w_i$ is a weighted average of the input vectors. With this notation, the solution of Eq. (4) is $\hat{r} = \sum_{i=1}^{L+K} \mu_i w_i^*$, where $\mathbf{w}^* = [w_1^*, \dots, w_{L+K}^*]$ satisfies the CE condition $F(\mathbf{w}^*) = G(\mathbf{w}^*)$.

This is the core problem solved by Plug and Play (PnP) algorithms [21]. In these algorithms, one or more of the proximal maps H_k are replaced with advanced denoisers, such as BM3D [22] or deep neural networks. In [17], it is shown that the solution \mathbf{w}^* can be found by Mann iterations of the operator $T = (2G - I)(2F - I)$. For a fixed parameter $\rho \in (0, 1)$, this takes the form $\mathbf{w}^{n+1} = (1 - \rho)\mathbf{w}^n + \rho T(\mathbf{w}^n)$.

The agents H_k should be chosen to enforce spatial regularity and together act as an algorithmic encoding of the prior information. In this work, we choose H_1 to be the application of BM3D on each xy -slice of the 3D volume, while H_2 and H_3 perform the same action on the yz and xz -slices, respectively. This approach, called multi-slice fusion [23,24], yields an efficient 3D volume regularizer.

The forward agents, the proximal maps $F_\ell(w; \sigma^2)$, are solved as follows. Omitting the subscript ℓ for clarity, we use the likelihood $p(y|r)$, and remove terms independent of r , to get the proximal map

$$F(w; \sigma^2) = \underset{r \in \mathbb{R}^N}{\operatorname{argmin}} \left\{ \log |AD(r)A^H + \sigma_\eta^2 I| + y^H (AD(r)A^H + \sigma_\eta^2 I)^{-1} y + \frac{1}{2\sigma^2} \|r - w\|^2 \right\}. \quad (8)$$

For computational efficiency, we assume A is orthogonal, with $A^H A = \alpha I$. In which case, the proximal map becomes

$$F(w; \sigma^2) \approx \underset{r_i \geq 0}{\operatorname{argmin}} \left\{ \sum_{i=1}^N \left(\log |r_i + \sigma_\eta^2 / \alpha| + \frac{|(A^H y)_i|^2}{r_i + \sigma_\eta^2 / \alpha} \right) + \frac{1}{\sigma^2} \|r - w\|^2 \right\}. \quad (9)$$

Since the cost function in Eq. (9) is separable in r_i , we can solve for each r_i independently of the others. The solution of this optimization problem can be shown to be a non-negative root of a cubic equation.

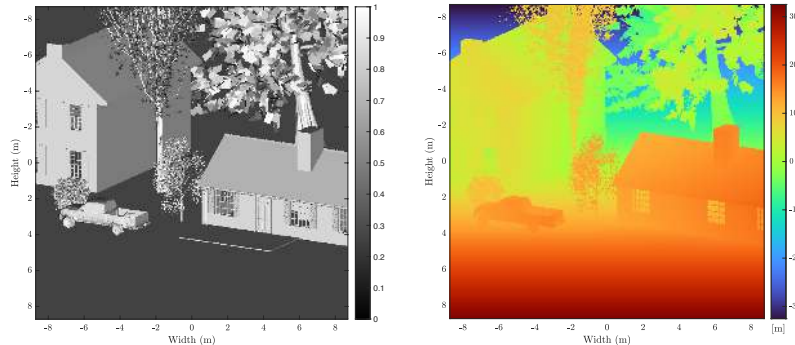


Figure 1. Target of simulated RCH detection.

4. Results

In this section, we present results from a set of simulated data. The subject of this experiment is shown in Figure 1. We present 3D images as a pair of intensity (left) and depth (right) maps. These are taken as the maximum intensity, and the location of the maximum intensity, along the imaging axis, respectively. We simulated the RCH detection of statistically independent speckle realizations of this scene at various signal-to-noise-ratios (SNRs). We compare the results of a traditional speckle-averaged reconstruction, $\sum_{\ell=1}^L |A^H y_{\ell}|^2 / L$, with the output of our proposed reconstruction algorithm. For each instance in the dataset, we initialize our algorithm with the speckle-averaged image and compute a maximum of 15 Mann iterations with $\rho = 0.5$. For each of the forward agents, we use a proximal map weight of $\sigma^2 = 0.01$. For the prior agents H_k , each MACE weights μ_i are $1/6$, with the remainder being evenly distributed amongst the L forward agents. For these prior agents, the strength of the denoiser is determined by a parameter within BM3D which we denote σ_B^2 . We vary this parameter according to the amount of noise in the data.

Assuming an RCH system operating at the shot noise limit, the SNR is proportional to the mean number of signal photoelectrons (pe) incident on the recording focal-plane array over the entire laser bandwidth. We, however, characterize the SNR in terms of the average number of photoelectrons per cross-range pixel of the reconstructed image, which we denote \bar{m} . These quantities are related by a scalar roughly equivalent to the number of pixels per diffraction-limited spot [26]. In our simulation, this quantity is approximately 7.6. Similarly, the volumetric noise variance σ_{η}^2 is related by a scalar roughly equivalent to the number of voxels in the main lobe of the PSF h , which is approximately 18 in our simulation. In summary, simulated detection of \bar{m} mean photoelectrons per pixel is equivalent to an SNR of approximately $7.6 * \bar{m}$ and a noise variance $\sigma_{\eta}^2 \approx 1 / (18 * \bar{m})$ over the diffraction-limited spot.

The values of \bar{m} were chosen to be $1/2, 1, 2, 4,$ and 8 mean photoelectrons per pixel with a corresponding denoiser strength, σ_B^2 , of $0.05, 0.04, 0.03, 0.025$ and 0.02 . Each of the five noise levels were obtained using $1, 2, 4,$ and 8 independent speckle realizations for a total of twenty imaging scenarios. Other parameters used in data simulation are summarized in Table 1.

Table 1. Data Simulation Parameters

Symbol	Parameter	Value
L_x, L_y, L_z	Target Extent	(17.5 m, 17.5 m, 65.0 m)
$\Delta x, \Delta y, \Delta z$	Voxel Size	(0.068 m, 0.068 m, 0.254 m)
dx, dy, dz	Image Resolution	(0.166 m, 0.166 m, 0.620 m)
D	Aperture Diameter	2 cm
λ	Central Wavelength	1550 nm
Z	Path Length	5 km
B	Bandwidth	0.239 GHz
$\Delta\omega$	Frequency Interval	2.30 MHz




Figure 2. Speckle-averaged and MACE reconstructions at various SNR.

Figures 2 and 3 show the MACE reconstruction (bottom) for select scenarios in the simulated data set. The results are compared to the standard speckle-averaged intensity reconstruction $\sum_{\ell=1}^L |A^H y_{\ell}|^2 / L$ (top). Each figure shows the reconstruction as a pair of intensity (left) and depth (right) maps. Figure 2 shows the results consisting of eight speckle realizations over various values of SNR, while Figure 3 shows the results for a fixed total SNR of 2 mean photoelectrons per pixel over a variable number of

speckle realizations. Through the combination of modeling the speckle-free surface reflectivity and 3D denoising priors, our method reconstructs smooth intensity values at accurate depths at significantly lower SNR than standard speckle-averaging. In cases with high speckle variation, our approach struggles to regularize the intensity values but is still able to remove volumetric background noise to find the surface and obtain accurate depth information.




Figure 3. Speckle-averaged and MACE reconstructions with variable number of looks.

In order to quantitatively analyze our algorithm, we use a metric based on a point cloud representation of the image. From a 3D image, we generate a point cloud by taking the maximum intensity in range for each cross-range pixel. From this point cloud, we compute the RMSE of the Euclidean distance from each point in our reconstruction to the closest point in a high-resolution reference point cloud, as done in [27]. This can be interpreted as a generalization of the RMSE on depth images as it allows for inter-pixel registration of points. Further, it allows the reconstruction to be compared to a high-resolution ground truth image without interpolation. One downside of this method is the heavy detriment outliers can have on the computation. However, these points can easily be removed by thresholding their distance to the ground truth, and their existence provides another metric for quality of reconstruction.

The point cloud RMSE for the speckle-averaged image (left) and MACE reconstruction (right) are shown in Figure 4. The MACE reconstruction sees an improvement in all cases. For example, the MACE reconstruction achieves similar metrics at 1 mean photoelectron per pixel as the speckle-averaged reconstruction does at 8 mean photoelectrons per pixel. Metrics for outlier removal show similar improvements. For high SNR, MACE approaches the minimum achievable point cloud RMSE.




Figure 4. Point cloud RMSE (m) of speckle-averaged (left) and MACE (right) reconstructions.

5. Conclusion

We present a new approach for RCH image reconstruction that couples advanced prior models with stochastic physics-based models. By first reformulating the inversion in terms of the real-valued reflectance, and further by utilizing the MACE framework to incorporate a BM3D image prior and data from multiple looks, our proposed method reduces the effects of measurement noise and speckle variation. Our results on simulated data indicate that incorporating computational imaging algorithms into RCH imaging systems may significantly reduce the laser power required to operate them.

6. Acknowledgements

This effort was supported in part by the U.S. Air Force through contract number FA8650-20-F-1925. The views expressed in this article are those of the authors and do not reflect on the official policy of the Air Force, Department of Defense, or the U.S. Government. C. A. Bouman was also supported in part by the Showalter Trust, and G. T. Buzzard was supported in part by NSF grant CCF-1763896.

7. References

- [1] Samuel T. Thurman and James R. Fienup. “Phase-Error Correction in Digital Holography”. In: *Journal of the Optical Society of America A* 25.4 (Apr. 2008), p. 983.
- [2] Samuel T. Thurman. “Phase-Error Correction in Digital Holography Using Single-Shot Data”. In: *JOSA A* 36.12 (Dec. 2019), pp. D47–D61.
- [3] Abbie E. Tippie and James R. Fienup. “Phase-Error Correction for Multiple Planes Using a Sharpness Metric”. In: *Optics Letters* 34.5 (Mar. 2009), p. 701.
- [4] Abbie E. Tippie and James R. Fienup. “Multiple-Plane Anisoplanatic Phase Correction in a Laboratory Digital Holography Experiment”. In: *Optics Letters* 35.19 (Oct. 2010), p. 3291.
- [5] Sennan Sulaiman, Steve Gibson, and Mark Spencer. “Predictive Dynamic Digital Holography and Image Sharpening”. In: *Journal of the Optical Society of America A* 35.6 (June 2018), p. 923.
- [6] C. J. Pellizzari, M. F. Spencer, and C. A. Bouman. “Coherent Plug-and-Play: Digital Holographic Imaging Through Atmospheric Turbulence Using Model-Based Iterative Reconstruction and Convolutional Neural Networks”. In: *IEEE Transactions on Computational Imaging* 6 (2020).
- [7] Casey J. Pellizzari, Mark F. Spencer, and Charles A. Bouman. “Imaging through Distributed-Volume Aberrations Using Single-Shot Digital Holography”. In: *JOSA A* 36.2 (Feb. 2019),
- [8] Casey J. Pellizzari and Charles A. Bouman. “Inverse Synthetic Aperture LADAR Image Construction: An Inverse Model-Based Approach”. In: *SPIE Optical Engineering + Applications*. (2016).
- [9] Casey J. Pellizzari et al. “Optically Coherent Image Formation and Denoising Using a Plug and Play Inversion Framework”. In: *Applied Optics* 56.16 (June 2017), p. 4735.
- [10] Casey J. Pellizzari et al. “Demonstration of Single-Shot Digital Holography Using a Bayesian Framework”. In: *Journal of the Optical Society of America A* 35.1 (Jan. 2018), p. 103.
- [11] N. H. Farhat. “Holography, Wave-length Diversity and Inverse Scattering”. In: *Optics in Four Dimensions-1980*. Ensenada, Mexico, 1980, pp. 627–642.
- [12] Joseph C. Marron and Timothy J. Schulz. “Three-Dimensional, Fine-Resolution Imaging Using Laser Frequency Diversity”. In: *Optics Letters* 17.4 (Feb. 1992), p. 285.
- [13] Louis Cuellar et al. “Digital Holography Three-Dimensional Imaging Using Frequency Chirping of a Laser”. In: *Unconventional Imaging and Adaptive Optics 2020*. Vol. 11508.
- [14] Jason W. Stafford, Bradley D. Duncan, and David J. Rabb. “Phase Gradient Algorithm Method for Three-Dimensional Holographic Ladar Imaging”. In: *Applied Optics* 55.17 (June 2016), p. 4611.
- [15] Wesley E. Farriss et al. “Sharpness-Based Correction Methods in Holographic Aperture Ladar (HAL)”. In: *Unconventional and Indirect Imaging, Image Reconstruction, and Wavefront Sensing 2018*.
- [16] Matthias T. Banet et al. “3D Multi-Plane Sharpness Metric Maximization with Variable Corrective Phase Screens”. In: *Applied Optics* 60.25 (Sept. 2021), G243.
- [17] Gregory T. Buzzard et al. “Plug-and-Play Unplugged: Optimization-Free Reconstruction Using Consensus Equilibrium”. In: *SIAM Journal on Imaging Sciences* 11.3 (Jan. 2018), pp. 2001–2020.
- [18] Joseph Goodman. *Introduction to Fourier Optics*, 4th Ed. Macmillan Learning, New York (2017).
- [19] Wesley Farriss. “Iterative Phase Estimation Algorithms in Interferometric Systems”. 2021. University of Rochester, Ph.D. thesis.
- [20] Vladimir V. Protopopov. “Principles of Optical Heterodyning”. In: *Laser Heterodyning*. Ed. by Vladimir V. Protopopov. Springer Series in Optical Sciences. Berlin, Heidelberg: Springer, 2009, pp. 1–49.
- [21] Singanallur V. Venkatakrishnan, Charles A. Bouman, and Brendt Wohlberg. “Plug-and-Play Priors for Model Based Reconstruction”. In: *2013 IEEE GlobalSIP*. Austin, TX: IEEE, Dec. 2013, pp. 945–948.
- [22] Kostadin Dabov et al. “Image Denoising with Block-Matching and 3D Filtering”. In: *Electronic Imaging 2006*. Ed. by Edward R. Dougherty et al. San Jose, CA, Feb. 2006.
- [23] Soumendu Majee et al. “4D X-Ray CT Reconstruction Using Multi-Slice Fusion”. In: *arXiv:1906.06601*
- [24] Maliha Hossain et al. “Ultra-Sparse View Reconstruction for Flash X-Ray Imaging Using Consensus Equilibrium”. In: *arXiv:2103.15979 [eess]* (Apr. 2021). *arXiv: 2103.15979 [eess]*.
- [25] Venkatesh Sridhar et al. “Fast Algorithms for Model-based Imaging through Turbulence”. In: *Proc. SPIE 11543, Artificial Intelligence and Machine Learning in Defense Applications II*.
- [26] M. F. Spencer. “Spatial Heterodyne”. In: *Encyclopedia of Modern Optics*. Vol. 4.
- [27] Thomas Welsh. “Quantitative Analysis of 3D Images Formed Using Range Compressed Holography”. 2017. University of Dayton, Master's thesis.