

WaveMo: Learning Wavefront Modulations to See Through Scattering

Mingyang Xie^{1*} Haiyun Guo^{2*} Brandon Y. Feng³ Lingbo Jin²
 Ashok Veeraraghavan² Christopher A. Metzler¹

¹University of Maryland

²Rice University

³Massachusetts Institute of Technology

Abstract

Imaging through scattering media is a fundamental and pervasive challenge in fields ranging from medical diagnostics to astronomy. A promising strategy to overcome this challenge is wavefront modulation, which induces measurement diversity during image acquisition. Despite its importance, designing optimal wavefront modulations to image through scattering remains under-explored. This paper introduces a novel learning-based framework to address the gap. Our approach jointly optimizes wavefront modulations and a computationally lightweight feedforward “proxy” reconstruction network. This network is trained to recover scenes obscured by scattering, using measurements that are modified by these modulations. The learned modulations produced by our framework generalize effectively to unseen scattering scenarios and exhibit remarkable versatility. During deployment, the learned modulations can be decoupled from the proxy network to augment other more computationally expensive restoration algorithms. Through extensive experiments, we demonstrate our approach significantly advances the state of the art in imaging through scattering media. Our project webpage is at <https://wavemo-2024.github.io/>.

1. Introduction

Imaging through scattering media presents a significant challenge across diverse scenarios, ranging from navigating with fog [2, 28, 30], rain [9, 48, 49], or murky water to recovering intricate structures through human skin and tissue [10, 42, 47]. The core of this challenge are the irregular phase delays light experiences as it scatters. These phase delays blur and warp any images captured through the scattering media. Effectively addressing this issue is crucial for unlocking new computer vision capabilities in fields such as medical imaging and astronomy.

From a frequency domain perspective, the effect of a

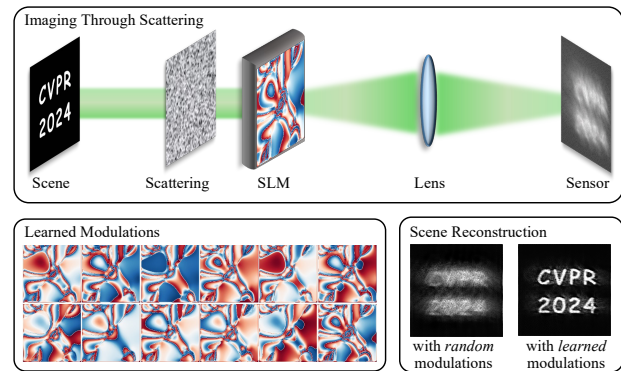


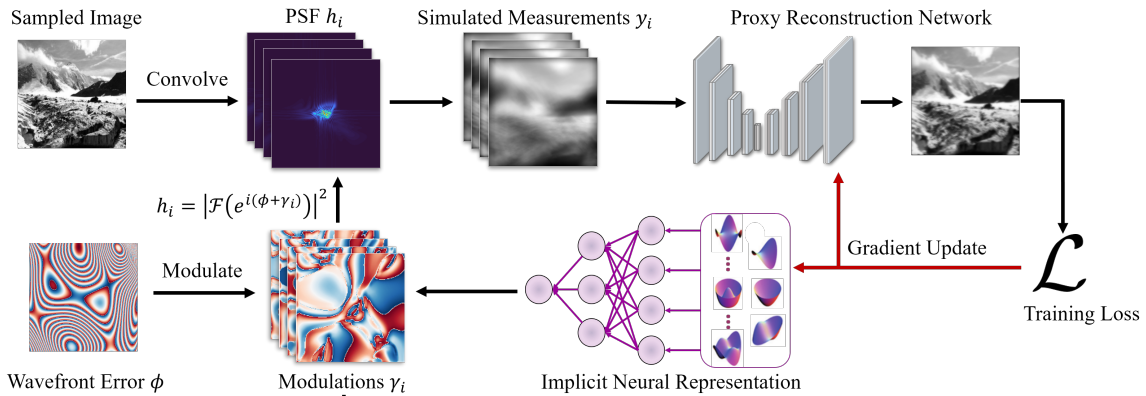
Figure 1. **Learned Wavefront Modulations.** *Top:* During acquisition, we can modulate the wavefront of scattered light by using a spatial light modulator (SLM), and capture a set of image measurements useful for scene reconstruction. *Bottom left:* We propose learning modulations that enhance our ability to recover the scattered scene. *Bottom right:* Our learned modulations drastically improve the reconstruction quality of a state-of-the-art method [8] that previously applies randomly chosen modulations.

scene passing through scattering is the destruction or filtering of frequency content. Recovering these lost frequencies is very difficult using purely computational methods, which are effectively forced to speculate on the missing frequencies based on those that have been preserved. A question naturally arises: Can we modify the measurement process to prevent the loss of information during acquisition?

A promising solution is to actively modulate the scattered wavefront during acquisition, in hopes that applying multiple modulations increases the chance that a frequency component is preserved. This approach is related to phase diversity imaging [11, 15, 17, 25, 27, 37, 41, 45], which typically captures multiple scattered images of a static scene with different modulations applied using deformable mirrors or spatial light modulators (SLM). However, phase diversity imaging traditionally was limited to imaging simple scenes with mild optical aberrations, and the phase modulations were often selected based on simple heuristics. The recent development of neural wavefront shaping (NeuWS) [8]

*Equal Contribution.

Training



Testing




Figure 2. **Overview of the Proposed End-to-end Learning Framework.** During training, we jointly optimize the implicit neural representation for the wavefront modulations and the proxy reconstruction network in an end-to-end fashion. During inference, one can apply the learned wavefront modulations to any reconstruction algorithms for imaging through scattering, either a trained feed-forward reconstruction network or an unsupervised iterative optimization algorithm [8]. The former approach is far faster and performs better when test data and training data fall into similar distributions, while the latter generalizes better to unseen distributions of target scenes.

brings a significant breakthrough by leveraging recent innovations in machine learning and enables imaging high-resolution and dynamic scenes through severe aberrations. Still, the wavefront modulations applied by NeuWS during acquisition are chosen randomly without explicit consideration of their impact on frequency preservation.

In this paper, we take a significant step towards a more systematic and principled approach in designing effective modulations that better preserve frequencies against scattering during image acquisition. Moving beyond the traditional reliance on simple heuristics for determining modulation patterns, our method combines the strengths of differentiable optimization and data-driven learning. Central to our approach is the novel integration of the optical model of wavefront modulations with a *proxy* reconstruction network, resulting in a fully differentiable system. This integration allows for the simultaneous optimization of both the proxy reconstruction network and the modulation patterns, which are optimized end-to-end based on a large, simulated training dataset. While imaging through scattering often requires iterative optimization algorithms because data-driven feedforward networks often fail to generalize outside their training domain, this paper demonstrates the remarkable finding that training a *hard-to-generalize* network can

become a useful proxy allowing us to learn *generalizable* modulations. Crucially, this proxy feedforward network allows us to bypass the expensive backpropagation computation through an iterative reconstruction algorithm, which would be necessary, but hopelessly impractical, if we were to naively optimize the modulations: each iteration would require us to finish an entire iterative reconstruction algorithm. As will be shown in our simulated and real data experiments, our learned modulations significantly enhance the reconstruction capability of both our data-driven proxy network and a state-of-the-art iterative optimization-based method [8] for imaging through scattering.

Our contributions are:

- We present a novel end-to-end learning framework to optimize acquisition-time wavefront modulations to enhance our abilities to see through scattering.
- We demonstrate in both simulated and real experiments that our learned modulations substantially improve image reconstruction quality and effectively generalize to unseen targets and scattering media.
- We show that the learned modulations can be decoupled from the jointly-trained proxy reconstruction network and significantly enhance the reconstruction quality of state-of-the-art unsupervised approaches.

2. Related Work

Learning-based Imaging Through Scattering Media.

A detailed review of imaging through scattering media can be found in [10]. We highlight a few recent learning-based approaches. Various data-driven approaches can be applied to image through scattering [5, 13, 20, 22, 35, 36]. They generally capture or synthesize large quantities of training data and then learn a mapping from corrupted measurements to reconstructed scenes. However, they tend to struggle with out-of-distribution data. There have also been unsupervised algorithms by fitting network-based parameterizations of the scene to a collection of measurements [8, 15], which can generalize across different distributions of aberrations and target scenes. Both of these methods take advantage of phase diversity.

End-to-End Learning of Optical Systems. End-to-end learning is a flexible learning-based system design framework that describes optical systems using parameterized differentiable optical models which can then be optimized with training data [31]. End-to-end learning has been applied successfully for extended depth-of-field imaging [14, 21, 29, 31], depth estimation [4, 14, 29, 43], high-dynamic-range imaging [23, 33], seeing through near-lens occlusions [30], and many other applications. To facilitate end-to-end learning, prior works have also incorporated the usage of a proxy network to mimic the response of a specific optical system or black-box algorithm [26, 39].

Phase Diversity Imaging. Phase diversity imaging (PDI) can recover a clear image of an aberrated scene by obtaining images with deliberately introduced known aberrations. PDI does not pose strong assumptions on the target [45] or the aberration [16], and the optics and calibration process required are relatively simple compared to other adaptive optics imaging [11]. While early applications of PDI are often constrained to simple aberrations due to computation complexities [32], the introduction of coded diffraction patterns has enhanced its capacity for reconstructing intricate scenes [3]. Recent PDI developments include improved wavefront sensing [45, 50], phase retrieval [27, 44], adaptive optics [7, 15], and imaging through scattering [6, 8, 46].

3. Learned Wavefront Modulations

Here we describe our approach for learning wavefront modulations through training a proxy reconstruction network.

3.1. Scattering Problem

This paper focuses on the scenario where capturing an image through scattering media can be modeled as

$$y = h(\phi) * x + \epsilon, \quad (1)$$

where y is the captured measurement, x is the target scene that we aim to reconstruct, ϵ is noise, and $h(\phi)$ is the unknown, spatially invariant point spread function (PSF) describing the optical scattering, which manifests as unknown phase delays ϕ to the wavefront. Assuming the target object is illuminated with spatially incoherent monochromatic light, the PSF h is related to the phase error ϕ by $h(\phi) = |\mathcal{F}(m \circ e^{j(\phi)})|^2$, where \mathcal{F} is the 2D Fourier transform, \circ is the hadamard product and m is the aperture mask [12]. Our goal is to recover x from y captured through ϕ .

3.2. Phase Diversity

Phase diversity imaging seeks additional information about the phase error ϕ by capturing a sequence of measurements y_1, y_2, \dots, y_K , each with additional known phase errors $\gamma_1, \gamma_2, \dots, \gamma_K$. As a result of these additional modulations, the imaging forward model becomes $y_i = h(\phi + \gamma_i) * x + \epsilon_i$. In traditional approaches, the phase modulations $\{\gamma_i\}_{i=1}^K$ were typically either simple defocus sweeps or randomly chosen patterns generated from Zernike polynomials [8, 25], and their reconstruction objective seeks to minimize

$$\mathcal{L}_{PDI}(\hat{x}, \phi) = \sum_{i=1}^K \|y_i - h(\phi + \gamma_i) * \hat{x}\|^2, \quad (2)$$

over \hat{x} and ϕ , where \hat{x} is the estimate of scene x , assuming ϵ_i follows an i.i.d. additive Gaussian distribution.

3.3. Learned Modulations

While previous works have leveraged differentiable proxies of physical systems to design better algorithms [26, 39], this work leverages a differentiable proxy, \mathcal{P} , of non-differentiable algorithms to design better physical systems—specifically the phase modulation $\Gamma = \{\gamma_i\}_{i=1}^K$ used to image through scattering.

We sample a training image x_n and simulate passing it through an optical aberration drawn from some known distribution while we apply K learnable modulations to the scattered wavefront. This model allows us to simulate a collection of K images $Y_n^\Gamma = \{y_1^n, y_2^n, \dots, y_K^n\}$, which are different scattered observations of x_n based on Eq. (1).

We then optimize both the proxy reconstruction algorithm \mathcal{P} and the modulation patterns Γ to minimize

$$\mathcal{L}_{learned}(\mathcal{P}, \Gamma) = \sum_{n=1}^N \|\mathcal{P}(Y_n^\Gamma) - x_n\|^2. \quad (3)$$

Note that the ultimate goal of the learning is not to design a set of modulations specific to the network \mathcal{P} . Rather, we use the performance of \mathcal{P} as a proxy to probe how effective the modulations Γ are for the scattering problem. The differentiability of \mathcal{P} allows for back-propagation to guide

the optimization of the modulations. As we will demonstrate, while \mathcal{P} itself may not generalize outside of its training data domain, the learned modulations Γ are compatible with other more generalizable reconstruction algorithms.

Our end-to-end training pipeline is illustrated in Figure 2. For the proxy reconstruction network, we choose a U-Net with skip connections and self-attention modules [24]. To regularize the highly non-convex problem of optimizing modulations, we use an implicit neural representation for the modulations Γ . Specifically, the MLP G takes a fixed vector Z with 28 channels, which corresponds to the first 28 Zernike polynomials [18]; it outputs a 16-channel vector corresponding to 16 modulation patterns: $\Gamma_G = G(Z)$.

Incorporating the implicit neural representation for the modulations, our final loss function becomes

$$\mathcal{L}_{final}(\mathcal{P}, G) = \sum_{n=1}^N \|\mathcal{P}(Y_n^{\Gamma_G}) - x_n\|^2, \quad (4)$$

where $Y_n^{\Gamma_G} = \{h(\phi_n + G(Z)_i) * x_n\}_{i=1}^K$.

3.4. Applying Learned Modulations

While the trained proxy network \mathcal{P} demonstrates considerable capability in recovering scenes from scattering in our experiments, supervised data-driven methods often fail to generalize to novel or cross-domain scenes, particularly when compared to unsupervised approaches.

Therefore, after obtaining the learned modulations Γ_G and applying them during real-world acquisition, we send the resulting modulated measurements to an unsupervised iterative optimization-based reconstruction algorithm [8] which does not suffer from the generalization issue of data-driven methods. In effect, our method offers a best-of-both-worlds scenario: the modulations Γ_G are effectively learned thanks to the joint supervised training with the proxy network \mathcal{P} , but their enhanced data acquisition quality is transferable to other reconstruction algorithms, thus allowing us to benefit from the generalization and domain adaptability inherent in unsupervised reconstruction methods.

4. Analysis of Learned Modulations

Before we show the performance of the learned modulations with real-world experiments in Section 6, we first analyze the effectiveness of learned modulations in simulation.

4.1. Impact on Frequency

When the scattering media scrambles the incoming wavefront, it destroys some of the target scene’s frequencies, either by destructively interfering with them or by scattering them outside the imaging sensor. If those frequencies are not captured in the measurement, it becomes difficult to faithfully recover them. By applying multiple learned wavefront modulation patterns, we increase the chance that each




Figure 3. **MTF Comparison.** We compare the MTF of learned wavefront modulations vs. the MTF of modulations randomly drawn from the same distribution of the scattering media. The X-axis represents spatial frequency, and the Y-axis represents the modulation transfer (the higher the better). Our learned wavefront modulations exhibit a higher MTF compared to unoptimized ones, especially over higher frequency bands, suggesting the learned modulations preserve more high-frequency information.

frequency is better preserved by at least one of the modulations. With the extra information from all modulated measurements, a better image reconstruction performance could be achieved regardless of the reconstruction algorithm.

Thus, we assess the performance of our learned modulations by their combined ability to preserve frequencies. We choose to the commonly used Modulation Transfer Function (MTF) as our metric [1, 19, 34, 38, 40]. MTF measures the contrast produced by an imaging system at different spatial frequencies, defined as $\text{MTF}_i^\omega(\phi) = |\mathcal{F}[h_i(\phi)]|^\omega$, where h_i is the PSF for the i -th modulation and the superscript ω denotes indexing the MTF at a specific frequency ω . For the multiple modulations in our problem setup, we define the combined MTF as:

$$\text{MTF}_{\text{comb}}^\omega(\{\gamma_i\}, \phi) = \max_i \text{MTF}_i^\omega(\phi + \gamma_i), \quad (5)$$

where for each frequency ω , we take the maximum over the MTFs calculated from all modulations.

In Figure 3, we plot the combined MTF for both the learned wavefront modulations and the random ones as baseline. The learned ones have a clear advantage at high frequencies (the higher the frequency is, the further away it is to the center of the X-axis), suggesting that the learned modulations are better at preserving finer details of the target scene. Note that during training, we never explicitly encourage the MTF to be higher — the only loss term we use is the mean square error between reconstruction and ground truth image. This implies that our end-to-end training implicitly encourages the modulations to preserve higher-frequency information.

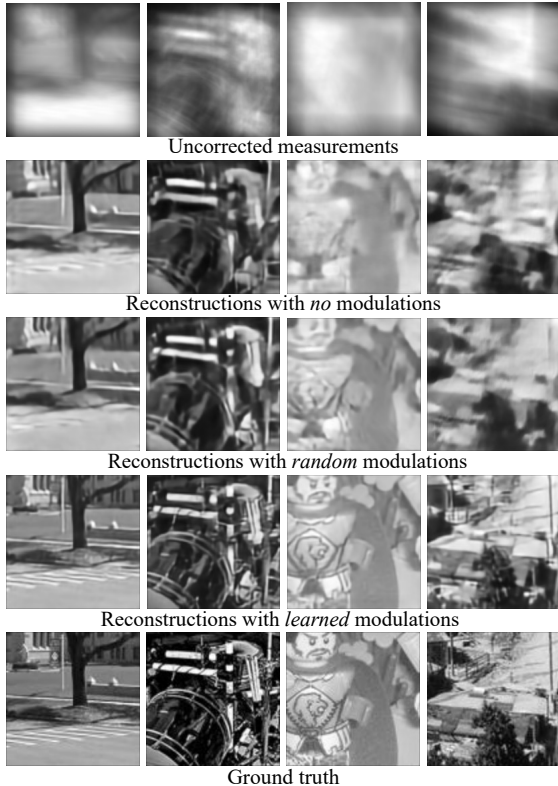


Figure 4. **Proxy Network Reconstruction on Simulated Scattering.** Simulated imaging results through different aberrations. Learned modulations lead to a better quality.

4.2. Impact on Reconstruction

We evaluate our jointly learned modulations and proxy network both quantitatively and qualitatively. Both the training and test set are from the MIT Places 365 Dataset [51]. Training details are provided later in Section 5. Here, we use two baselines: (1) simply training a feed-forward reconstruction network without any wavefront modulations, and (2) training a feed-forward reconstruction network with randomly chosen, fixed wavefront modulations as used in prior PDI work. For a fair comparison, our approach and baselines use the same network architecture and are trained on the same data for the same number of iterations. Their only difference is the usage of wavefront modulations.

As shown in Figure 4, compared to baselines, reconstructions using learned modulations demonstrate finer details and higher consistency with the ground truth. Table 1 shows the quantitative comparison averaged over a test set of 1,000 images, where our approach outperforms the baselines by over 3.9 dB. While the results above are all from the jointly trained proxy reconstruction network, we will show in Section 6 that our learned wavefront modulations can decouple from the proxy network and also improve reconstruction quality based on iterative optimization [8].

Metric	Modulations		
	None	Random	Learned
PSNR	25.476	26.439	30.391
SSIM	0.7640	0.7980	0.9082

Table 1. **PSNR and SSIM of Simulation Results.** We report the test-time performance of our jointly trained modulations and the proxy network. We compare it against two baselines where either we do not use wavefront modulations or use randomly chosen, unoptimized modulations. The metrics are averaged over the reconstruction of 1000 test images, each of which is scattered with a randomly sampled optic aberration that is unseen during training. The results demonstrate the huge performance improvement from the optimization of wavefront modulations.

5. Implementation

Software Implementation & Training. For the proxy reconstruction network, we use an attention U-Net [24]; for the MLP that represents the 16 wavefront modulations, we use a two-layer MLP with leaky-ReLU activation. In each training iteration, the optical aberrations are sampled randomly from Zernike basis functions. The standard deviation of the coefficients for each Zernike basis is randomly generated from a uniform distribution from 5 to 6. We train our proposed framework for 2 million iterations on the MIT Places Dataset [51] with the Adam Optimizer and a learning rate of 0.001. Training took 12 hours on an NVIDIA RTX A6000 GPU. For experimental results, we finetune our trained model on the captured bio-tissue data under the same training settings except with a learning rate of 0.0001.

Hardware Implementation. Our optical configuration is depicted in Figure 6. A continuous-wave laser with a wavelength of 532 nm was collimated and subsequently passed through a laser speckle reducer (Optotune LSR-4C-L). The resulting spatially incoherent light was then polarized and reflected by a digital micromirror device (DMD). The DMD (LightCrafter™ 4500) display is capable of rapidly presenting 8-bit target images. For ease of assembly, a reflecting mirror was incorporated after the DMD. Then, the light is focused by lens L0 ($f = 100\text{mm}$) onto the front focal plane of L1, a $10\times$ microscopic objective. Our optical aberration—resin painted on glass—was placed at the same location. L1 and a tube lens L2 ($f = 200\text{mm}$) constitute a $4f$ system, and thereby the SLM is manipulating the Fourier plane. We used a HOLOEYE LETO-3 phase-only SLM to modulate the wavefront. Finally, the light reflected by the SLM was imaged on the camera (1384×1036 pixel Grasshopper3) via L3 ($f = 500\text{mm}$). Our capture speed is limited to 30 fps by the camera frame rate.

Experimental Dataset. Our experimental data are from human pathological slides, imaged with optical microscopy (Zeiss Microscope Axio Imager.A2) and a $10\times$ objective.




Figure 5. **Proxy Network Reconstruction on Physical Scattering.** Experimental results of imaging objects through scattering media by using our proxy reconstruction network. The left columns are in-distribution adipose tissue slides (zoomed-in region labeled with red boxes); the right columns are out-of-distribution targets. Learned modulations yield superior imaging quality for both in-distribution and out-of-distribution scenes, with the former out-performing the latter.




Figure 6. **Optical Setup.** A spatially incoherent light source illuminates the DMD and passes through thick scattering. The DMD displays the target images for training and testing. The SLM is placed onto the Fourier plane and projects learned patterns. The wavefront is modulated and subsequently imaged onto a camera.

Each image is 256×256 pixels. Our dataset includes 1,000 sectional images of adipose tissues. 40 of these images were utilized for testing and the rest were used for training. To evaluate the generalizability across different human tissues, we also imaged sections of human stomach and glandular epithelium, trachea, and finger. Additionally, we tested with targets outside the domain of the training data, such as numbers and letters.

6. Experimental Results

We employed the optical system depicted in Figure 6. Our learned modulations significantly enhance the imaging capabilities through scattering media. We tested on datasets from both similar and distinctly different distributions from our experimental training data. Our learned modulations consistently yield superior visual reconstruction results. Furthermore, we evaluated these modulations on both static and dynamic scenes. We show that our learned modulations significantly enhance reconstruction in both cases, which suggests a broad applicability of our approach.

6.1. Proxy Reconstruction Network

The proxy network, fine-tuned on 960 adipose tissue images through random scattering, was evaluated on a test set from the same tissue type, as shown in the left of Figure 5. Reconstructions driven by measurements of 16 learned SLM patterns preserved the original contrast and high-frequency details of the dyed sections. The close-ups show remarkable detail accuracy. By comparison, reconstructions trained on measurements of random or no modulation struggled, particularly in lower-intensity areas, yielding blurry outputs. Beyond achieving good reconstruction fidelity on objects similar to our training dataset, we extended our evaluation to vastly different data distributions, like numbers and letters. The results, as shown on the right of Figure 5, clearly indicate that the learned modulations yield bet-

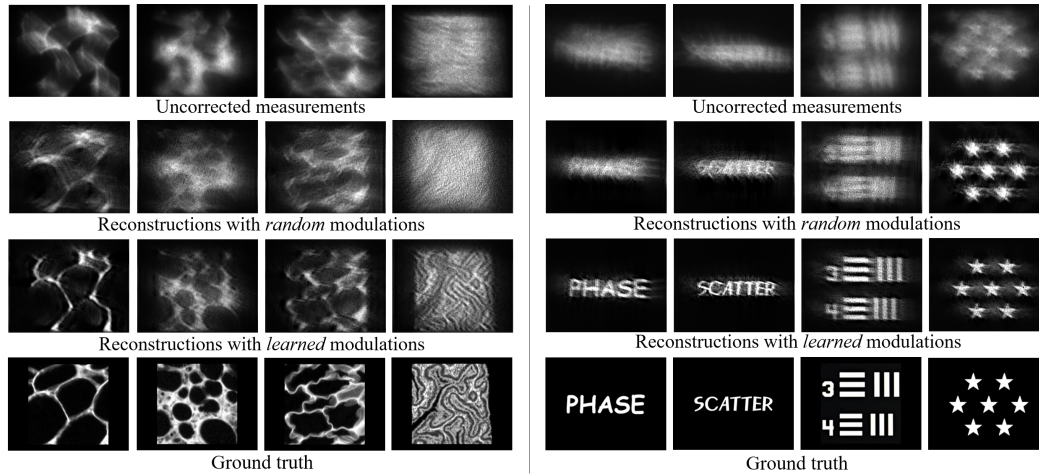


Figure 7. **Unsupervised Reconstruction on Physical Scattering with Static Scenes.** Experimental results of imaging different static targets through scattering media using learned wavefront modulations with unsupervised iterative optimization [8]. Learned modulations enhance reconstruction quality.

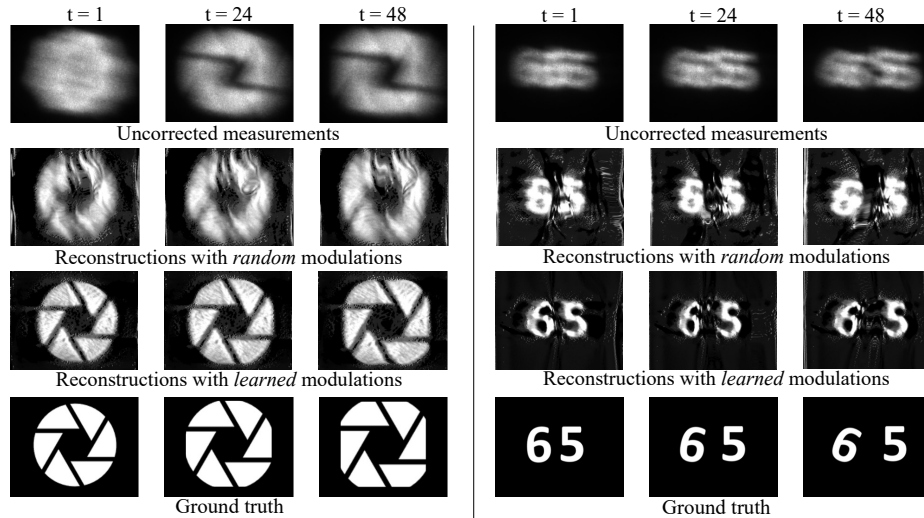


Figure 8. **Unsupervised Reconstruction on Physical Scattering with Dynamic Scenes.** Experimental results of imaging an opening camera aperture through severe aberrations (left) and imaging non-regional motions (number 6 rotating clockwise and number 5 translating to the right) through severe aberrations. We use the unsupervised iterative reconstruction approach [8]. Learned modulations highly improve the quality of the reconstructed dynamic scene.

ter reconstruction quality, particularly in the USAF target and stars. PSNR and SSIM are computed in Tables 2 and 3. Among the table cells, “None” means that the measurements for each target only comprised the unmodulated image (no modulation patterns are applied); “Random” means that the modulation patterns are randomly sampled from the same distribution as the aberrations.

6.2. Unsupervised Iterative Approaches

While the proxy network \mathcal{P} exhibits considerable effectiveness in real-world experiments, its performance remains unsatisfactory for out-of-distribution data, which is expected. Nevertheless, our learned modulations are designed to further augment generalizable, untrained reconstruction algorithms based on iterative optimizations. To validate this strategy, we deploy our learned modulations to a recently developed unsupervised iterative approach for imaging through scattering, NeuWS [8], in our benchmark study.

Method	Data type	Modulation		
		None	Random	Learned
Proxy	Tissue	16.53	17.57	19.06
Proxy	Out-of-dist.	9.34	9.90	10.71
Iterative [8]	Static	N/A	11.26	14.61
Iterative [8]	Dynamic	N/A	8.90	12.89

Table 2. **PSNR of Experimental Results.** For our jointly trained feed-forward proxy reconstruction network (“proxy”), we tested on 40 tissue samples and 8 out-of-distribution scenes, all of which are static. For the iterative method [8], we tested on the same 8 out-of-distribution static scenes. We also tested [8] on 2 dynamic scenes, each with 48 frames. The iterative method relies on multiple wavefront modulations and therefore cannot recover objects with a single measurement, hence the “N/A”. Compared against random modulations or no modulations, our learned modulations lead to better reconstruction performance for both the proxy network and the unsupervised iterative approach.

Method	Data type	Modulations		
		None	Random	Learned
Proxy	Tissue	0.44	0.48	0.58
Proxy	Out-of-dist.	0.29	0.30	0.32
Iterative [8]	Static	N/A	0.21	0.38
Iterative [8]	Dynamic	N/A	0.23	0.33

Table 3. **SSIM of Experimental Results.** Same as Table 2 but showing SSIM. Our learned modulations performs better.

Figure 7 shows the reconstruction performance of NeuWS [8] on 8 static objects, each using only 16 measurements. The superiority of using learned modulations is visually evident in the results. We excluded experiments with no modulations here, due to the algorithm’s inability to recover objects using a single unmodulated measurement.

For dynamic scenes, each scene involves 48 frames captured by cycling through the 16 SLM patterns. Figure 8 demonstrates the reconstruction results on dynamic scenes: the result on the left illustrates a camera aperture gradually enlarging; the result on the right shows two different motions to solve the non-regional dynamic reconstruction (the number 6 on the left rotates clockwise at 0.5° per frame, while the number 5 translates from the center to the right). Quantitative metrics on the dynamic results are included in Table 2 and 3. Full video results are in the [project webpage](#).

7. Discussion & Conclusion

Role of the Proxy Reconstruction Network. The proxy network not only aims to reconstruct the images; its performance also indicates how good the wavefront modulations are for the restoration task. Therefore, joint differentiable optimization of both the proxy network and the modulations works in our favor, and the experimental results indeed in-

dicating that this is an effective way of learning useful wavefront modulations. Without this proxy network, we would have to somehow connect the modulation update with the performance of an iterative optimization method, which is impractical due to time and memory constraints.

Regularization Perspective. To recover missing signal frequencies due to scattering, most existing methods rely on prior information of the scene, either hand-crafted or data-driven. This work rethinks the strategy of incorporating prior knowledge: *can we impose a prior over wavefront modulations?* The goal here is to let the resulting imaging system capture more information about the target scene through scattering. Our results validate the effectiveness of learning a prior over the wavefront modulation domain.

Applicability to Other Vision Tasks. The end-to-end learning approach can potentially benefit other vision tasks through scattering media, using a task-specific proxy network. For instance, if we were to do object detection or semantic segmentation through scattering media, we can replace the proxy reconstruction U-Net with a network tailored for these two tasks, and perform end-to-end training in the same way as before. Note that we validated our proposed method on a monochromatic imaging system, which is prevalent in medical and scientific imaging. Extending it to broadband imaging tasks, e.g., outdoor navigation, requires the optimization of wavelength-dependent modulation patterns, which we leave for future work.

Conclusion. We proposed a robust end-to-end learning framework integrating the optical scattering model with a proxy image reconstruction network. The learned wavefront modulations can both work with the proxy reconstruction network and augment a generalizable, optimization-based algorithm. We conducted intensive experiments to validate our approach. This work shows the synergy of advanced wavefront modulation techniques with cutting-edge machine learning methods, representing a significant leap forward in computer vision and optical imaging.

Acknowledgements

This work was supported in part by AFOSR Young Investigator Program award no. FA9550-22-1-0208, ONR award no. N00014-23-1-2752, seed grant from SAAB, Inc., a seed grant from the UMD Brain and Behavior Institute, NSF Expeditions award no. CCF-1730574, ONR Scattering media award no. N00014-23-1-2714, and NSF PaThs up award no. EEC-1648451. We thank Sachin Shah for helpful discussions and Kevin F. Kelly for providing the DMD.

References

- [1] Jesse K Adams, Dong Yan, Jimin Wu, Vivek Boominathan, Sibao Gao, Alex V Rodriguez, Soonyoung Kim, Jennifer Carns, Rebecca Richards-Kortum, Caleb Kemere, et al. In vivo lensless microscopy via a phase mask generating diffraction patterns with high-contrast contours. *Nature Biomedical Engineering*, 6(5):617–628, 2022. 4
- [2] Codruta Orniana Ancuti, Cosmin Ancuti, and Radu Timofte. Nh-haze: An image dehazing benchmark with non-homogeneous hazy and haze-free images. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1798–1805, 2020. 1
- [3] Emmanuel J Candes, Xiaodong Li, and Mahdi Soltanolkotabi. Phase Retrieval Via Wirtinger Flow: Theory and Algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015. 3
- [4] Julie Chang and Gordon Wetzstein. Deep optics for monocular depth estimation and 3d object detection. In *Proc. IEEE ICCV*, 2019. 3
- [5] Dongdong Chen, Mike Davies, Matthias J Ehrhardt, Carola-Bibiane Schönlieb, Ferdia Sherry, and Julián Tachella. Imaging with equivariant deep learning: From unrolled network design to fully unsupervised learning. *IEEE Signal Processing Magazine*, 40(1):134–147, 2023. 3
- [6] Wei-Yu Chen, Matthew O’Toole, Aswin C Sankaranarayanan, and Anat Levin. Enhancing Speckle Statistics for Imaging Inside Scattering Media. *Optica*, 9(12):1408–1416, 2022. 3
- [7] Santiago Echeverri-Chacón, René Restrepo, Carlos Cuartas-Vélez, and Néstor Uribe-Patarroyo. Vortex-enhanced coherent-illumination phase diversity for phase retrieval in coherent imaging systems. *Optics letters*, 41(8):1817–1820, 2016. 3
- [8] Brandon Y Feng, Haiyun Guo, Mingyang Xie, Vivek Boominathan, Manoj K Sharma, Ashok Veeraraghavan, and Christopher A Metzler. Neuws: Neural wavefront shaping for guidestar-free imaging through static and dynamic scattering media. *Science Advances*, 9(26):eadg4671, 2023. 1, 2, 3, 4, 5, 7, 8
- [9] Xueyang Fu, Jia-Bin Huang, Delu Zeng, Yue Huang, Xinghao Ding, and John William Paisley. Removing rain from single images via a deep detail network. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1715–1723, 2017. 1
- [10] Sylvain Gigan, Ori Katz, Hilton B de Aguiar, Esben Ravn Andresen, Alexandre Aubry, Jacopo Bertolotti, Emmanuel Bossy, Dorian Bouchet, Joshua Brake, Sophie Brasselet, et al. Roadmap on Wavefront Shaping and Deep Imaging in Complex Media. *Journal of Physics: Photonics*, 4(4):042501, 2022. 1, 3
- [11] Robert A Gonsalves. Phase retrieval and diversity in adaptive optics. *Optical Engineering*, 21(5):829–832, 1982. 1, 3
- [12] Joseph W. Goodman. *Introduction to Fourier Optics*. Roberts & Co., Englewood, Colo, 3rd ed. edition, 2005. 3
- [13] Xiaowen Hu, Jian Zhao, Jose Enrique Antonio-Lopez, Stefan Gausmann, Rodrigo Amezcua Correa, and Axel Schülzgen. Adaptive inverse mapping: a model-free semi-supervised learning approach towards robust imaging through dynamic scattering media. *Optics Express*, 31(9):14343–14357, 2023. 3
- [14] Hayato Ikoma, Cindy M. Nguyen, Christopher A. Metzler, Yifan Peng, and Gordon Wetzstein. Depth from defocus with learned optics for imaging and occlusion-aware depth estimation. *IEEE International Conference on Computational Photography (ICCP)*, 2021. 3
- [15] Iksung Kang, Qinrong Zhang, Stella X Yu, and Na Ji. Coordinate-based neural representations for computational adaptive optics in widefield microscopy. *arXiv preprint arXiv:2307.03812*, 2023. 1, 3
- [16] Richard L Kendrick, Daniel S Acton, and AL Duncan. Phase-diversity wave-front sensor for imaging systems. *Applied Optics*, 33(27):6533–6546, 1994. 3
- [17] Peter Kner. Phase diversity for three-dimensional imaging. *JOSA A*, 30(10):1980–1987, 2013. 1
- [18] Vasudevan Lakshminarayanan and Andre Fleck. Zernike polynomials: a guide. *Journal of Modern Optics*, 58(7):545–561, 2011. 4
- [19] Juhyun Lee, Jinsoo Jeong, Jaebum Cho, Dongheon Yoo, Byoung-hyo Lee, and Byoung-ho Lee. Deep neural network for multi-depth hologram generation and its training strategy. *Optics Express*, 28(18):27137–27154, 2020. 4
- [20] Yunzhe Li, Yujia Xue, and Lei Tian. Deep speckle correlation: a deep learning approach toward scalable imaging through scattering media. *Optica*, 5(10):1181–1190, 2018. 3
- [21] Xin Liu, Linpei Li, Xu Liu, Xiang Hao, and Yifan Peng. Investigating deep optics model representation in affecting resolved all-in-focus image quality and depth estimation fidelity. *Opt. Express*, 30(20):36973–36984, 2022. 3
- [22] Christopher A Metzler, Felix Heide, Prasana Rangarajan, Muralidhar Madabhushi Balaji, Aparna Viswanath, Ashok Veeraraghavan, and Richard G Baraniuk. Deep-Inverse Correlography: Towards Real-Time High-Resolution Non-Line-of-Sight Imaging. *Optica*, 7(1):63–71, 2020. 3
- [23] Christopher A Metzler, Hayato Ikoma, Yifan Peng, and Gordon Wetzstein. Deep optics for single-shot high-dynamic-range imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1375–1385, 2020. 3
- [24] Ozan Oktay, Jo Schlemper, Loïc Le Folgoc, M. J. Lee, Mattias P. Heinrich, Kazunari Misawa, Kensaku Mori, Steven G. McDonagh, Nils Y. Hammerla, Bernhard Kainz, Ben Glocker, and Daniel Rueckert. Attention u-net: Learning where to look for the pancreas. *ArXiv*, abs/1804.03999, 2018. 4, 5
- [25] Richard G Paxman, Timothy J Schulz, and James R Fienup. Joint estimation of object and aberrations by using phase diversity. *JOSA A*, 9(7):1072–1085, 1992. 1, 3
- [26] Yifan Peng, Suyeon Choi, Nitish Padmanaban, and Gordon Wetzstein. Neural Holography With Camera-in-the-Loop Training. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020. 3
- [27] Nikolaj Reiser, Min Guo, Hari Shroff, and Patrick J La Riviere. Phase diverse phase retrieval for microscopy: Com-

- parison of gaussian and poisson approaches. *arXiv preprint arXiv:2308.00734*, 2023. 1, 3
- [28] Wenqi Ren, Sibao Liu, Hua Zhang, Jin shan Pan, Xiaochun Cao, and Ming-Hsuan Yang. Single image dehazing via multi-scale convolutional neural networks. In *European Conference on Computer Vision*, 2016. 1
- [29] Sachin Shah, Sakshum Kulshrestha, and Christopher A Metzler. Tidy-psfs: Computational imaging with time-averaged dynamic point-spread-functions. *ICCV*, 2023. 3
- [30] Zheng Shi, Yuval Bahat, Seung-Hwan Baek, Qiang Fu, Hadi Amata, Xiao Li, Praneeth Chakravarthula, Wolfgang Heidrich, and Felix Heide. Seeing through obstructions with diffractive cloaking. *ACM Transactions on Graphics (TOG)*, 41(4):1–15, 2022. 1, 3
- [31] Vincent Sitzmann, Steven Diamond, Yifan Peng, Xiong Dun, Stephen Boyd, Wolfgang Heidrich, Felix Heide, and Gordon Wetzstein. End-to-end optimization of optics and image processing for achromatic extended depth of field and super-resolution imaging. *ACM Trans. Graph.*, 37(4), 2018. 3
- [32] Carlas S Smith, Arnold J den Dekker, Raluca Andrei, Rufus Fraanje, and Michel Verhaegen. Fast phase diversity wavefront sensing using object independent metrics. In *Adaptive Optics Systems III*, pages 2250–2258. SPIE, 2012. 3
- [33] Qilin Sun, Ethan Tseng, Qiang Fu, Wolfgang Heidrich, and Felix Heide. Learning rank-1 diffractive optics for single-shot high dynamic range imaging. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1386–1396, 2020. 3
- [34] Qilin Sun, Congli Wang, Fu Qiang, Dun Xiong, and Heidrich Wolfgang. End-to-end complex lens design with differentiable ray tracing. *ACM Trans. Graph.*, 40(4):1–13, 2021. 4
- [35] Yu Sun, Zhihao Xia, and Ulugbek S Kamilov. Efficient and Accurate Inversion of Multiple Scattering with Deep Learning. *Optics Express*, 26(11):14678–14688, 2018. 3
- [36] Waleed Tahir, Hao Wang, and Lei Tian. Adaptive 3d descattering with a dynamic synthesis network. *Light: Science & Applications*, 11(1):42, 2022. 3
- [37] Brian J Thelen, Richard G Paxman, David A Carcarara, and John H Seldin. Maximum a posteriori estimation of fixed aberrations, dynamic aberrations, and the object from phase-diverse speckle data. *JOSA A*, 16(5):1016–1025, 1999. 1
- [38] Lei Tian and Laura Waller. Quantitative differential phase contrast imaging in an led array microscope. *Optics express*, 23(9):11394–11403, 2015. 4
- [39] Ethan Tseng, Felix Yu, Yuting Yang, Fahim Mannan, Karl St. Arnaud, Derek Nowrouzezahrai, Jean-Francois Lalonde, and Felix Heide. Hyperparameter optimization in black-box image processing using differentiable proxies. *ACM Transactions on Graphics (TOG)*, 38(4), 2019. 3
- [40] Hakan Urey, Ned Nestorovic, Baldwin S Ng, and Abraham A Gross. Optics designs and system mtf for laser scanning displays. In *Helmet-and Head-Mounted Displays IV*, pages 238–248. SPIE, 1999. 4
- [41] Curtis R Vogel, Tony F Chan, and Robert J Plemmons. Fast algorithms for phase-diversity-based blind deconvolution. In *Adaptive Optical System Technologies*, pages 994–1005. SPIE, 1998. 1
- [42] Abbie T Watnik and Dennis F Gardner. Wavefront Sensing in Deep Turbulence. *Optics and Photonics News*, 29(10):38–45, 2018. 1
- [43] Yicheng Wu, Vivek Boominathan, Huaijin Chen, Aswin Sankaranarayanan, and Ashok Veeraraghavan. Phasecam3d — learning phase masks for passive single view depth estimation. In *2019 IEEE International Conference on Computational Photography (ICCP)*, pages 1–12, 2019. 3
- [44] Meng Xiang, An Pan, Jinpeng Liu, Teli Xi, Xin Guo, Fei Liu, and Xiaopeng Shao. Phase diversity-based fourier ptychography for varying aberration correction. *Frontiers in Physics*, 10:129, 2022. 3
- [45] Qi Xin, Guohao Ju, Chunyue Zhang, and Shuyan Xu. Object-independent image-based wavefront sensing approach using phase diversity images and deep learning. *Optics express*, 27(18):26102–26119, 2019. 1, 3
- [46] Tomer Yeminy and Ori Katz. Guidestar-Free Image-Guided Wavefront Shaping. *Science Advances*, 7(21), 2021. 3
- [47] Seokchan Yoon, Moonseok Kim, Mooseok Jang, Youngwoon Choi, Wonjun Choi, Sungsam Kang, and Wonshik Choi. Deep Optical Imaging Within Complex Scattering Media. *Nature Reviews Physics*, 2(3):141–158, 2020. 1
- [48] Yi Yu, Wenhan Yang, Yap-Peng Tan, and Alex C Kot. Towards robust rain removal against adversarial attacks: A comprehensive benchmark analysis and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6013–6022, 2022. 1
- [49] Fan Zhang, Shaodi You, Yu Li, and Ying Fu. Learning rain location prior for nighttime deraining. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13148–13157, 2023. 1
- [50] Peiguang Zhang, Chengliang Yang, Zihao Xu, Zhaoliang Cao, Quanquan Mu, and Li Xuan. High-accuracy wavefront sensing by phase diversity technique with bisymmetric defocus diversity phase. *Scientific Reports*, 7(1):15361, 2017. 3
- [51] Bolei Zhou, Àgata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40:1452–1464, 2018. 5