

Unrolled Primal-Dual Networks for Lensless Cameras

OLIVER KINGSHOTT, University College London, United Kingdom

NICK ANTIPA, University of California San Diego, United States of America

EMRAH BOSTAN, ams OSRAM, Switzerland

KAAN AKŞIT, University College London, United Kingdom

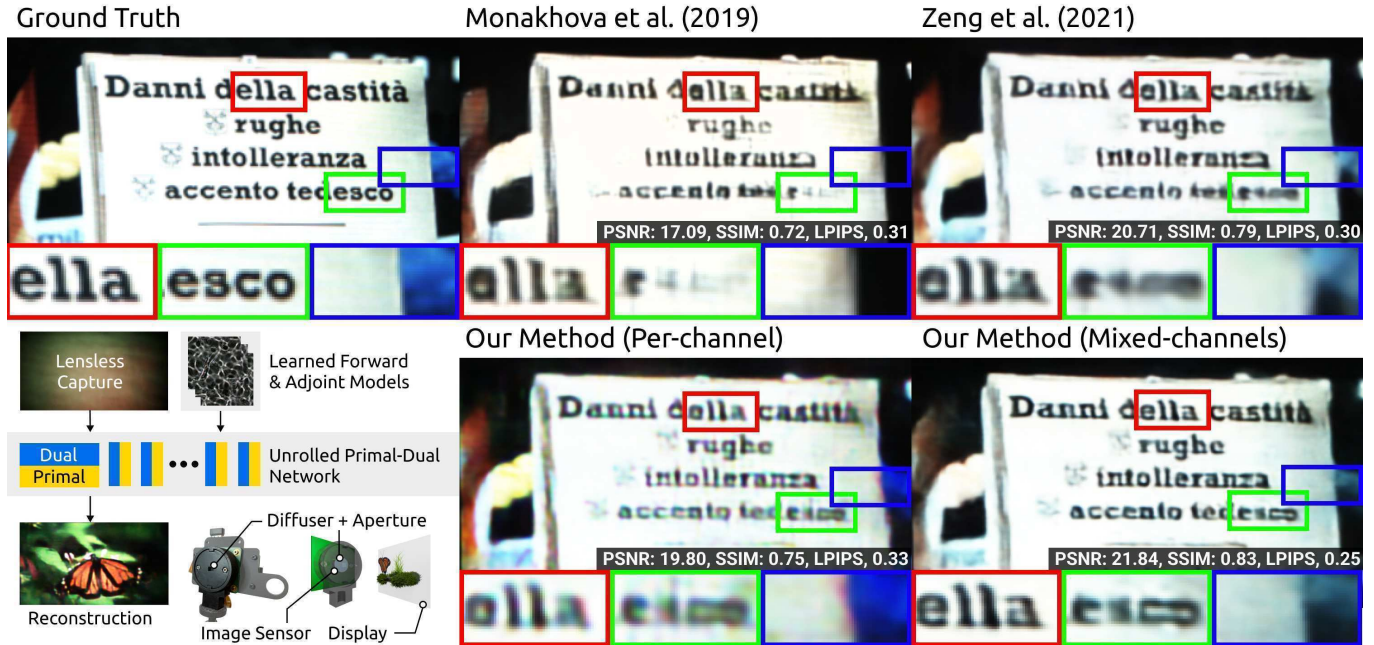


Fig. 1. Comparison of our unrolled primal-dual network with state of the art. Intensive post-processing of lensless images cannot correct the model error, over-smoothing images and removing important features, such as text. We propose to replace classical lensless reconstruction methods with our physically-informed unrolled primal-dual model, where the model includes a series of learned forward and adjoint models (pseudo point-spread functions and their inverse). As a result, our work can produce plausible images and recover additional features while reducing the need for deep post-processing networks such as U-Nets [Ronneberger et al. 2015] (Source image courtesy MIR Flickr [Huiskes and Lew 2008]).

Conventional image reconstruction models for lensless cameras often assume that each measurement results from convolving a given scene with a single experimentally measured point-spread function. These image reconstruction models fall short in simulating lensless cameras truthfully as these models are not sophisticated enough to account for optical aberrations or scenes with depth variations. Our work shows that learning a supervised primal-dual reconstruction method results in image quality matching state of the art in the literature without demanding a large network capacity. This improvement stems from our primary finding that embedding learnable forward and adjoint models in a learned primal-dual optimization framework can even improve the quality of reconstructed images (+5dB PSNR)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Woodstock '18, June 03–05, 2018, Woodstock, NY

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>

compared to works that do not correct for model error. In addition, we built a proof-of-concept lensless camera prototype that uses a pseudo-random phase mask to demonstrate our point. Finally, we share the extensive evaluation of our learned model based on an open dataset and a dataset from our proof-of-concept lensless camera prototype.

CCS Concepts: • **Computing methodologies** → **Computational photography**.

Additional Key Words and Phrases: lensless imaging, neural networks

ACM Reference Format:

Oliver Kingshott, Nick Antipa, Emrah Bostan, and Kaan Akşit. 2022. Unrolled Primal-Dual Networks for Lensless Cameras. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

A lensless camera uses a thin mask in place of a conventional lens. Masks can manipulate phase, amplitude, or the entire complex light field of a given scene. Unlike lenses in conventional cameras, these

masks can be placed near the imaging sensor, enabling thinner and lighter imaging systems. Additionally, lensless cameras offer the benefits of compressed imaging [Fergus et al. 2006; Liutkus et al. 2014], embedding higher dimensional scene information such as depth from a single capture. To benefit from these qualities, experts typically model lensless cameras as a linear system and recover images computationally by solving the inverse problem.

Pseudo-random phase masks have demonstrated adequate performance for lensless photography [Antipa et al. 2018; Boominathan et al. 2020]. Unfortunately, image reconstruction typically requires computationally expensive and slow iterative reconstruction algorithms (e.g. ADMM [Antipa et al. 2018] and FISTA [Beck and Teboulle 2009]). To address this, a growing number of works use data-driven Convolutional Neural Networks (CNNs) to improve the speed and quality of lensless image reconstructions [Bae et al. 2020; Barbastathis et al. 2019; Sinha et al. 2017]. A typical CNN with a limited receptive field size fails to accurately model the light transport of the imaging system [Goodman 2005], leading to learned models which fail to reconstruct lensless images accurately and efficiently. Recent literature proposes neural networks that include a physical model with a large receptive field [Boominathan et al. 2020; Monakhova et al. 2019]. These neural networks typically use a single-shot calibration measurement of the Point-Spread Function (PSF) to represent the physical model of the imaging system. However, without the use of precisely engineered masks [Boominathan et al. 2020; Tseng et al. 2021], image formation in lensless cameras cannot be fully expressed by a single PSF model [Yanny et al. 2020]. This model mismatch can lead data-driven regularizers to hallucinate missing features or create overly smooth images. Therefore, the development of models that can correct for model error without increased computational complexity or extensive calibration is of critical importance for the widespread adoption of lensless imaging. Our proposed method replaces ADMM with a learned optimization scheme, improving image quality by reducing model error as opposed to intensive post-processing. The result is a versatile deeply-calibrated lensless imaging architecture that avoids model error in the resulting reconstructions. We provide the results of numerous experiments comparing our method against existing image reconstruction algorithms for lensless cameras.

Specifically, our work provides the following contributions:

- Learned primal-dual for lensless imaging. We show for the first time that a modified learned primal-dual optimization framework [Adler and Öktem 2018] can recover images from a lensless camera using a pseudo-random phase mask.
- Learned forward-adjoint model. We embed additional linear operators within our learned primal-dual framework. These learned forward-adjoint models are jointly optimized with the rest of our model using the same paired training examples. We show that our extended model provides a significant visual quality enhancement in our image reconstructions. Our method promises reductions up to 50% in reconstruction error while using a fraction of the parameters compared to previous works.
- Lensless camera prototype. We build a proof of concept lensless camera to test further and demonstrate the performance

of our model in an actual lensless camera with a pseudo-random mask. We provide an automatic calibration routine that can train our model without the need for an additional camera with a conventional lens.

Limitations. When compared to models that use a single calibrated forward model, our method yields an improvement in the quality of lensless image reconstructions. However, a thorough investigation is required to identify explainable links between our learned forward models and physically accurate models in the future. In our experiments with our in-house built camera, we observe a lesser quality in image reconstructions when compared with the state of the art datasets [Boominathan et al. 2020; Monakhova et al. 2019]. We believe these originate from the fact that the off-the-shelf diffuser we use does not fully resemble the case that we draw our inspiration from [Antipa et al. 2018]. However, our work significantly improves the image quality both on benchmark datasets [Monakhova et al. 2019] and our in-house built camera.

2 RELATED WORK

We introduce a novel image reconstruction method for lensless cameras. Here, we provide a brief survey of prior art in lensless cameras, unsupervised lensless image reconstruction methods and learned image reconstruction techniques. Curious readers can read more about lensless cameras through the work by Boominathan et al. [Boominathan et al. 2022].

2.1 Lensless cameras

The idea of building cameras without requiring optical lenses has been a long-standing vision for scientists [Barker 1920] as optical lenses can be bulky, hard to manufacture with great precision, and are typically focused at one plane at a time. The advent of ubiquitous high performance computing and the promise of high dimensional capture has led to a resurgence of interest in lensless cameras. Mask based lensless cameras have been demonstrated with coded illumination [Zheng and Asif 2021], coded apertures [Asif et al. 2017; Horisaki et al. 2020], amplitude-only diffraction gratings (e.g., pinhole arrays [Anand et al. 2020]), photon sieves [Yöntem et al. 2018], separable amplitude masks [DeWeert and Farm 2015], Fresnel Zone plates [Wu et al. 2020]), phase-only diffraction gratings [Antipa et al. 2018; Bernet et al. 2011] and metalenses [Tseng et al. 2021]. Additionally, the mask used in a lensless imaging system can also be co-designed with an algorithm that recovers scene information [Tseng et al. 2021]. The depth-varying PSFs of phase mask imaging systems can augment existing 2D imaging sensors with near-field 3D imaging [Antipa et al. 2018]. Alternatively, single-pixel detectors combined with coded illumination patterns can be used for time-based imaging [Huang et al. 2013; Satat et al. 2017].

In our work, we show a lensless camera prototype for experimental validation. Our prototype is similar to the one demonstrated by Antipa et al. [2018] but differs in implementation details, which we go through in our implementation section.

2.2 Unsupervised Lensless Image Reconstruction Methods

The large spatial extent of the PSFs used in phase-mask based lensless cameras necessitates a cropped convolution model, owing to

the limited size of the imaging sensor. By modelling the convolution and the sensor crop as separable sub-problems, the Alternating-Direction Method of Multipliers [Antipa et al. 2018] can be used to recover images using convex optimization. However, modelling field-varying aberrations is cumbersome process using convex optimization approaches, typically requiring a 10x or greater increase in computational cost [Yanny et al. 2020].

2.3 Learned Lensless Image Reconstruction Methods

The advent of learning-based approaches eases the computational burden of lensless image reconstruction. The work by Monakhova et al. [2019] unrolls five iterations of ADMM and uses a large U-Net [Ronneberger et al. 2015] to improve perceptual quality. However, this approach has a limited ability to correct for model error in the resulting reconstructions, relying on intensive post-processing to achieve plausible reconstructed images. To our knowledge, the work by Rego et al. [2021] demonstrates the first attempt at implementing a blind deconvolution model for lensless cameras without involving PSF measurements. Our model requires re-training for each phase mask, yielding higher quality lensless reconstructions at the cost of portability. Khan et al. [2020] propose a fast learned reconstruction model for lensless cameras. By improving boundary conditions inherent in the sensor crop, they show that they can recover realistic images in a single step without the need for an iterative model. Our work embeds multiple large kernels within an unrolled iterative model to better compensate for optical aberrations. Zeng and Lam [2021] tackles model mismatch caused by imperfect modelling mainly due to spatially-varying PSFs with varying eccentricity. They achieve this by learning residual blocks during each unrolled iteration of ADMM, which are fed into the U-Net denoiser to correct for model error. We show that our method yields accurate intermediate reconstructions by separating the role of the denoising network from the model reconstruction network. Most recently, Yanny et al. [2022] have proposed pairing multiple Wiener filters with convolutional neural networks to recover accurate images in a lensless microscopy application. However, their method requires a experimental verification for phase-mask based lensless cameras as it targets microscopy.

In conclusion, existing learned methods depend both on accurate PSF calibration and additional training data to develop a suitable image prior. Our method makes better use of supervision by diverting trainable parameters towards improving the underlying physical model of light transport. By directly correcting for model-error, our method produces accurate intermediate reconstructions that are more consistent with images captured by a lensed camera. To our knowledge, our learned method delivers results that are on-par with the current state of the art in terms of speed and image quality, while offering greater parameter efficiency than previous works.

3 METHOD

We first introduce the forward model for a phase-mask based imaging system. We then present our proposed lensless image reconstruction model. Finally, we illustrate our deep calibration procedure which captures the necessary dataset for our supervised model-based reconstruction.

3.1 Problem Formulation

We assume that measurements from our imaging system, b , are the result of a linear transformation A applied to points in the scene x , with some additional noise ϵ :

$$b = Ax + \epsilon, \quad (1)$$

where b and x are vectors.

Each column of A corresponds to the linear transformation of a single point in the scene, also known as PSFs. Storing PSFs for each point in memory is a demanding task. Rather than storing all PSFs, using an aperture enables the approximation of A as a cropped convolution with a PSF measured along the optical axis [Antipa et al. 2018]

$$b = C(\text{PSF} * x) + \epsilon \quad (2)$$

Here, $*$ represents a circular convolution and C represents a crop down to the size of the imaging sensor. The lateral shifting of the large PSF outside of the bounds of the image sensor necessitates this cropped convolution model. A single experimentally measured PSF is typically used to reconstruct images using the described convolutional forward model [Antipa et al. 2018; Monakhova et al. 2019]. The on-axis PSF is typically measured by shining a point light source along the optical axis of an existing system. Under the assumption that b is the result of a cropped convolution with an experimentally measured PSF, we recover an estimate of the scene x by solving a regularized optimization problem:

$$\hat{x} \leftarrow \arg \min_x \frac{1}{2} \|C(\text{PSF} * x) - b\|_2^2 + \lambda \mathcal{R}(x), \quad (3)$$

where \mathcal{R} is a regularization function that penalizes unlikely solutions in the presence of noise, with λ controlling the amount of regularization with respect to the data fidelity term.

In this work, we seek to improve the quality of lensless imaging by embedding learnable convolution kernels that are the same size as the PSF within a learned optimization scheme.

3.2 Learning Large Kernels with Physically Informed Networks

Access to paired training examples unlocks a vast landscape of learned reconstruction techniques. Purely data-based architectures, such as U-Nets, typically require large numbers of paired training examples and suffer from poor generalization on unseen data. These limitations can be overcome by incorporating knowledge of physical processes, such as light transport [Kavaklı et al. 2022], into the neural network architecture. Physically informed networks such as learned primal-dual [Adler and Öktem 2018] are highly data-efficient, requiring only a moderate number of training examples, and tend to generalize well to unseen data. With access to paired training examples, but without knowledge of the true linear system A , we propose to train a reconstruction network \mathcal{G} with the goal of minimizing the average mean squared distance to ground truth reconstructions from a lensed camera x_{gt} :

$$\mathcal{L}_{MSE} := \|\mathcal{G}_\theta(b) - x_{\text{gt}}\|_2^2 \quad (4)$$

In the next section, we explain the design of \mathcal{G}_θ . As the focus of our work is to recover the signal encoded in b , we exclusively use mean-squared error as our loss function.

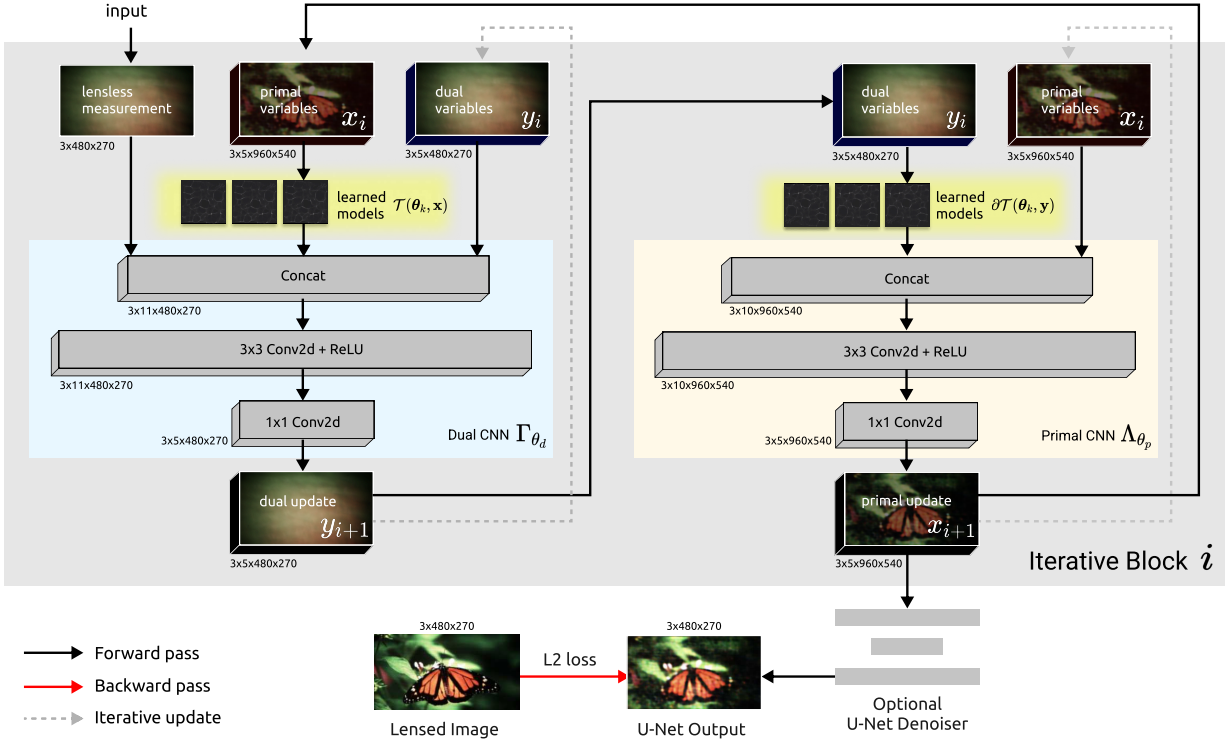


Fig. 2. Unrolled Primal-Dual network architecture for reconstructing lensless images. Our model accepts inputs in the form of a batch of RGB lensless measurements with a predetermined width and height. The blue box illustrates our dual update step, where variables in the measurement domain ($\{y_i, y_{i-1}, b\} \in Y$) are concatenated channel-wise before passing through two convolutional layers parameterized by θ_d^i . The yellow box illustrates our primal update step, where each variable in the measurement domain ($\{x_i, x_{i-1}\}$) is likewise concatenated and convolved with two layers parameterized by θ_p^i . Our forward-adjoint model tensor, θ_k , which is initialized with the value of a PSF measured using a point light source, is also optimized at each epoch. Finally, our trained model reconstructs images from lensless measurements.

3.2.1 Learned Primal Dual with a Physical Model. We propose a modified learned primal-dual architecture as our learned reconstruction network \mathcal{G} (Equation 4). Figure 2 illustrates how our data and parameters flow through the network. We extend the original work by Adler and Öktem [2018] in three ways. First, we replace the forward operator \mathcal{T} and its adjoint $\partial\mathcal{T}$ with the cropped convolution operation of our lensless camera in Equation (2):

$$\begin{aligned} \mathcal{T}(x) &\leftarrow C(\text{PSF} * x) \\ \partial\mathcal{T}(y) &\leftarrow P(\text{PSF} \star y), \end{aligned} \quad (5)$$

where P represents zero padding up to twice the size of the imaging sensor, and \star represents circular cross-correlation. $x \in X$ and $y \in Y$ are primal and dual variables respectively, with the former belonging to the domain of reconstructed images X and the latter in the domain of lensless measurements Y .

Second, we allow the PSF to be optimized during training. We initialize $\theta_k \leftarrow \text{PSF}$, allowing the network to modify the physical PSF during training:

$$\begin{aligned} \mathcal{T}(x) &\leftarrow C(\theta_k * x) \\ \partial\mathcal{T}(y) &\leftarrow P(\theta_k \star y), \end{aligned} \quad (6)$$

Finally, we wish to learn multiple kernels to improve our estimate of the true physical system. We choose to learn n convolution kernels, equal to the number of primal and dual variables. Let

$$\begin{aligned} x &= [x^1 \quad x^2 \quad \dots \quad x^n] \\ y &= [y^1 \quad y^2 \quad \dots \quad y^n], \\ \theta_k &= [\theta_k^1 \quad \theta_k^2 \quad \dots \quad \theta_k^n], \end{aligned} \quad (7)$$

then each primal and dual variable $x^{1\dots n}, y^{1\dots n}$ is convolved or cross-correlated with its own learned kernel $\theta_k^{1\dots n}$

$$\begin{aligned} \mathcal{T}(x) &\leftarrow C(\theta_k * x) \\ \partial\mathcal{T}(y) &\leftarrow P(\theta_k \star y). \end{aligned} \quad (8)$$

The above modifications result in a variation of the learned primal-dual algorithm with the following update steps:

$$\begin{aligned} y_i &\leftarrow \Gamma_{\theta_d^i}(y_{i-1}, \mathcal{T}(x_{i-1}), b) \\ x_i &\leftarrow \Lambda_{\theta_p^i}(x_{i-1}, \partial\mathcal{T}(y_i)), \end{aligned} \quad (9)$$

where $\Gamma_{\theta_d^i}, \Lambda_{\theta_p^i}$ are small convolutional neural networks that are parameterized by each unrolled iteration $i \in 1 \dots 10$. At the end

of the unrolled iterations, the variable x_{10}^1 is chosen as our best estimate of \hat{x} .

3.3 Per-channel & Mixed-channel models

To improve the performance of our method against baseline image quality metrics such as PSNR and SSIM, we propose an additional model based on higher dimensional feature maps as opposed to RGB images. Specifically, we replace k learned RGB kernels with $3 \times k$ single channel kernels, allowing for cross-channel communication across feature maps. This results in a model with an increased signal-to-noise performance at the cost of a decrease in subjective color accuracy. We provide a visual comparison of these two models and quantitative metrics in our results section.

4 IMPLEMENTATION

In this section we document the development of our own lensless camera as shown in Figure 1. Additional details are provided in the supplementary material.

Camera Design. We use a Raspberry Pi High-Quality camera connected to a Raspberry Pi Zero W. This specific camera features a removable lens housing which we replaced with our own 3D printed design. Following Monakhova et al. [2019], we used a 0.5 degree engineered diffuser as our mask, placed ~ 10 mm away from the image sensor. Our 3D printed housing is also illustrated in Figure 1. Our custom housing ensures that the optical element is placed at the desired distance from the imaging sensor, and contains space for an optional infrared filter.

Data Capture. To capture a training and test dataset, we place our camera ~ 15 cm away from a 5.5 inch OLED display. We illuminate a 5×5 square grid of pixels in the center of the display and capture the resulting image to measure the on-axis PSF. We then use FISTA [Beck and Teboulle 2009] to reconstruct a test image. This test image is used to estimate a homography that warps each ground truth image to match the perspective of the lensless camera. Automated software shows a variety of images from the DIV2K dataset [Timofte et al. 2018], capturing 8000 training images and 1000 test images.

5 EVALUATION

We first present the results of comparing our method against two central state-of-the-art work that uses DiffuserCam dataset [Monakhova et al. 2019; Zeng and Lam 2021]. We additionally perform ablation studies to determine the contribution from each component in our method on reconstructed image quality. Finally, we verify our method using our hardware prototype.

5.1 DiffuserCam results

We compare our model’s results against the work that uses DiffuserCam dataset [Monakhova et al. 2019] in Table 1, where the number of parameters used, the size of training and testing examples, processing time, and image quality are considered.

Our results suggest that our proposed method improves the quality of images reconstructed from measurements captured by a lensless camera. This is supported by qualitative results in Figure 3,

which appear to reproduce features that are more faithful to the original ground truth images.

5.2 Ablation Studies

Disabling U-Net Denoiser. To further confirm that the quality of our reconstructions has increased as a result of correcting for model error, we measure the quality of intermediate reconstructions without the use of a U-Net for denoising. We show our qualitative results in Figure 4 and quantitative results in Table 1. When our U-Net is disabled, the resulting images are noisy but are faithful to the ground truth images. Our intermediate reconstructions demonstrate that our model-based reconstruction network performs the bulk of the work in producing usable lensless reconstructions.

Effect of learning multiple models. We ran an additional study to quantify the effect of decreasing the number of learned models from 5 to 1. We include quantitative results in Table 1 and present reconstructed images from our reduced model in Figure 4. Decreasing the number of learned models from 5 to 1 decreases the resulting image quality after post-processing by ~ 2 dB.

5.3 Prototype results

We additionally compare the results of our learned model using a prototype camera built in the lab. We present sample reconstructions in Figure 5 and provide additional reconstructions in our supplementary material.

6 DISCUSSION

Comparison to classical methods. Our proposed models are end-to-end differentiable. They are trained to learn an unrolled iterative reconstruction algorithm, a physically informed model, and a suitable image prior. While our model appears to produce accurate intermediate reconstructions, it is difficult to discretely map each learned component of the model to a specific component existing classical methods. One line of future work could be to establish whether embedding learnable physical models within a classical variational method can achieve similar results. A forward model that is learned independently of image priors and a chosen reconstruction algorithm could be used to evaluate the data fidelity of reconstructed images against their lensless measurements.

Comparison to learned methods. When compared to learned methods that use a fixed PSF calibration measurement, our method is able to reconstruct images that more closely resemble images captured by a lensed camera. It is clear that the improved performance of our method is achieved by redistributing model parameters away from deep neural networks and towards the underlying physical model of light transport in lensless cameras. However, the exact mechanism through which our model improves performance against existing learned methods is unclear. It is possible that our model could be correcting for field-varying aberrations that are not captured by a single on-axis calibration measurement. However, we note that our proposed methods lack any explicit mechanism to apply each learned model to a specific spatial region. Finally, we note that our claim of improved data fidelity can only be measured implicitly by comparing our reconstructions with a lensed camera. In future work,




Fig. 3. Comparison of reconstructed test images against the ground truth images. We compare our method against MMCN [Zeng and Lam 2021]. MMCN is based on five unrolled iterations of ADMM with additional residual blocks to correct for model error. Our per-channel model (RGB) improves subjective color accuracy while our mixed-channel model (Mixed) recovers higher frequency content. The primary feature of both models is that multiple large kernels are learned to correct for model error.




Fig. 4. Comparison of our learned model-based reconstruction networks against unsupervised ADMM (converged). U-Net denoising was disabled to show that our intermediate reconstructions consist of images that are more faithful to the ground truth data. Learning additional kernels appears to improve accuracy while yielding results faster than classical methods. We reason that our network prioritises consistency with the true physical model, resulting in fewer artifacts.

Method	PSFs	U-Net	PSNR	LPIPS	Parameters	Runtime (ms)	Training Examples	Iterations
ADMM	1 RGB (fixed)		11.97	0.60	-	1190	0	100
Le-ADMM	1 RGB (fixed)		11.89	0.57	20	50	100	5
Le-ADMM-U	1 RGB (fixed)	✓	20.46	0.37	10.6M	55	24,000	5
Ours (RGB)	1 RGB (learned)		16.74	0.54	0.4M	74	9,000	10
		✓	21.47	0.43	1.2M	77	9,000	10
Ours (RGB)	5 RGB (learned)		16.91	0.51	2.0M	80	9,000	10
		✓	23.48	0.40	2.7M	88	9,000	10
Ours (Mixed)	15 (learned)		19.00	0.48	2.0M	82	9,000	10
		✓	25.34	0.35	3.8M	84	9,000	10

Table 1. Comparison of our models against previous work by Monakhova et al. [2019]. Our model achieves produces modestly accurate reconstructions quickly without the use of a large U-Net, at the cost of learning additional large kernels θ^k . These kernels occupy the majority of our parameter space. Adding a small U-Net to our models improves reconstruction quality further. Increasing the number of learned kernels improves PSNR by ~ 2 dB when combined with U-Net denoising, with cross-channel denoising adding another ~ 2 dB.



Fig. 5. Reconstructions from our lensless camera prototype trained using our RGB model. Our optical element consists of a thick holographic diffuser (0.76mm) with bulk scattering, leading to a degradation in image quality.

we would like to use measured or simulated field-varying PSFs to design robust models that can explicitly correct for field-varying aberrations without the need for manual calibration.

Color Accuracy. Our two proposed models highlight a potential trade-off between the recovery of high frequency details and color accuracy in phase mask cameras. Allowing the mixing of color channels appears to increase the frequency content of recovered images. However, our informal subjective opinion is that our per-channel model is able to reproduce color more accurately. We suspect that our per-channel model is vulnerable to color fringing artifacts introduced by the chosen phase masks. Future work could investigate treatment through the use of additional loss functions (such as those proposed by Heide et al. [2013]) or through improved phase mask design [Boominathan et al. 2020].

7 CONCLUSION

Unconventional camera designs with thin masks in place of conventional lenses offer freedom from the constraints of traditional optics. However, the speed of reconstruction and image quality in mask-based lensless camera designs remains a significant drawback. We argue that neural networks with embedded physical priors for lensless imaging can help to counter this drawback. We show that such an approach can provide on-par image reconstruction quality without demanding extensive resources in training. Thus, we hope that our work can further develop performant and interpretable methods for lensless image reconstruction.

ACKNOWLEDGEMENT

We thank Laura Waller, Kristina Monakhova, Tianjiao Zeng and Edmund Lam for their support in providing useful insights from their work; Tobias Ritschel for fruitful discussions at the early phases of the project; Koray Kavaklı for his support in hardware prototype related figure and camera homography related software; Tim Weyrich for dedicating GPU resource. Kaan Akşit and Oliver Kingshott relied on the Royal Society’s RGS\R2\212229 - Research Grants 2021 Round 2 for building the hardware prototype.

REFERENCES

- Jonas Adler and Ozan Öktem. 2018. Learned primal-dual reconstruction. *IEEE transactions on medical imaging* 37, 6 (2018), 1322–1332.
- Vijayakumar Anand, Soon Hock Ng, Jovan Maksimovic, Denver Linklater, Tomas Katkus, Elena P Ivanova, and Saulius Juodkazis. 2020. Single shot multispectral multidimensional imaging using chaotic waves. *Scientific reports* 10, 1 (2020), 1–13.
- Nick Antipa, Grace Kuo, Reinhard Heckel, Ben Mildenhall, Emrah Bostan, Ren Ng, and Laura Waller. 2018. DiffuserCam: lensless single-exposure 3D imaging. *Optica* 5 (1 2018), 1. Issue 1. <https://doi.org/10.1364/optica.5.000001>
- M. Salman Asif, Ali Ayremlou, Aswin Sankaranarayanan, Ashok Veeraraghavan, and Richard G. Baraniuk. 2017. FlatCam: Thin, Lensless Cameras Using Coded Aperture and Computation. *IEEE Transactions on Computational Imaging* 3, 3 (2017), 384–397. <https://doi.org/10.1109/TCL.2016.2593662>
- Donggeon Bae, Jaewoo Jung, Nakkyu Baek, and Seung Ah Lee. 2020. Lensless Imaging with an End-to-End Deep Neural Network. In *2020 IEEE International Conference on Consumer Electronics - Asia (ICCE-Asia)*. 1–5. <https://doi.org/10.1109/ICCE-Asia49877.2020.9276865>
- George Barbastathis, Aydogan Ozcan, and Guohai Situ. 2019. On the use of deep learning for computational imaging. *Optica* 6, 8 (Aug 2019), 921–943. <https://doi.org/10.1364/OPTICA.6.000921>

- HW Barker. 1920. Pin-hole or Lensless Photography. *Journal of the Royal Astronomical Society of Canada* 14 (1920), 16.
- Amir Beck and Marc Teboulle. 2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences* 2 (2009), 183–202. Issue 1. <https://doi.org/10.1137/080716542>
- Stefan Bernet, Walter Harm, Alexander Jesacher, and Monika Ritsch-Marte. 2011. Lensless digital holography with diffuse illumination through a pseudo-random phase mask. *Optics express* 19, 25 (2011), 25113–25124.
- Vivek Boominathan, Jesse K. Adams, Jacob T. Robinson, and Ashok Veeraraghavan. 2020. PhlatCam: Designed Phase-Mask Based Thin Lensless Camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42 (7 2020), 1618–1629. Issue 7. <https://doi.org/10.1109/TPAMI.2020.2987489>
- Vivek Boominathan, Jacob T Robinson, Laura Waller, and Ashok Veeraraghavan. 2022. Recent advances in lensless imaging. *Optica* 9, 1 (2022), 1–16.
- Michael J DeWeert and Brian P Farm. 2015. Lensless coded-aperture imaging with separable Doubly-Toeplitz masks. *Optical Engineering* 54, 2 (2015), 023102.
- Rob Fergus, Antonio Torralba, and William T Freeman. 2006. *Random lens imaging*. Technical Report. MIT.
- Joseph W Goodman. 2005. *Introduction to Fourier optics*. Roberts and Company Publishers.
- Felix Heide, Mushfiqur Rouf, Matthias B Hullin, Bjorn Labitzke, Wolfgang Heidrich, and Andreas Kolb. 2013. High-quality computational imaging through simple lenses. *ACM Transactions on Graphics (TOG)* 32, 5 (2013), 1–14.
- Ryoichi Horisaki, Yuka Okamoto, and Jun Tanida. 2020. Deeply coded aperture for lensless imaging. *Optics Letters* 45, 11 (2020), 3131–3134.
- Gang Huang, Hong Jiang, Kim Matthews, and Paul Wilford. 2013. Lensless imaging by compressive sensing. In *2013 IEEE International Conference on Image Processing*. 2101–2105. <https://doi.org/10.1109/ICIP.2013.6738433>
- Mark J Huiskes and Michael S Lew. 2008. The mir flicker retrieval evaluation. In *Proceedings of the 1st ACM international conference on Multimedia information retrieval*. 39–43.
- Koray Kavaklı, Hakan Urey, and Kaan Aksit. 2022. Learned holographic light transport. *Applied Optics* 61, 5 (2022), B50–B55.
- Salman Siddique Khan, Varun Sundar, Vivek Boominathan, Ashok Veeraraghavan, and Kaushik Mitra. 2020. FlatNet: Towards Photorealistic Scene Reconstruction from Lensless Measurements. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020), 1–1. <https://doi.org/10.1109/tpami.2020.3033882>
- Antoine Liutkus, David Martina, Sébastien Popoff, Gilles Chardon, Ori Katz, Geoffroy Lerosey, Sylvain Gigan, Laurent Daudet, and Igor Carron. 2014. Imaging with nature: Compressive imaging using a multiply scattering medium. *Scientific reports* 4, 1 (2014), 1–7.
- Kristina Monakhova, Joshua Yurtsever, Grace Kuo, Nick Antipa, Kyrollos Yanny, and Laura Waller. 2019. Learned reconstructions for practical mask-based lensless imaging. *Optics Express* 27 (9 2019), 28075. Issue 20. <https://doi.org/10.1364/oe.27.028075>
- Joshua D Rego, Karthik Kulkarni, and Suren Jayasuriya. 2021. Robust Lensless Image Reconstruction via PSF Estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 403–412.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 234–241.
- Guy Satat, Matthew Tancik, and Ramesh Raskar. 2017. Lensless Imaging With Compressive Ultrafast Sensing. *IEEE Transactions on Computational Imaging* 3, 3 (2017), 398–407. <https://doi.org/10.1109/TCI.2017.2684624>
- Ayan Sinha, Justin Lee, Shuai Li, and George Barbastathis. 2017. Lensless computational imaging through deep learning. *Optica* 4, 9 (Sep 2017), 1117–1125. <https://doi.org/10.1364/OPTICA.4.001117>
- Radu Timofte, Shuhang Gu, Jiqing Wu, Luc Van Gool, Lei Zhang, Ming-Hsuan Yang, Muhammad Haris, et al. 2018. NTIRE 2018 Challenge on Single Image Super-Resolution: Methods and Results. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Ethan Tseng, Shane Colburn, James Whitehead, Luocheng Huang, Seung-Hwan Baek, Arka Majumdar, and Felix Heide. 2021. Neural Nano-Optics for High-quality Thin Lens Imaging. *arXiv preprint arXiv:2102.11579* (2021).
- Jiachen Wu, Hua Zhang, Wenhui Zhang, Guofan Jin, Liangcai Cao, and George Barbastathis. 2020. Single-shot lensless imaging with fresnel zone aperture and incoherent illumination. *Light: Science & Applications* 9, 1 (2020), 1–11.
- Kyrollos Yanny, Nick Antipa, William Liberti, Sam Dehaeck, Kristina Monakhova, Fanglin Linda Liu, Konlin Shen, Ren Ng, and Laura Waller. 2020. Miniscope3D: optimized single-shot miniature 3D fluorescence microscopy. *Light: Science & Applications* 9, 1 (2020), 1–13.
- Kyrollos Yanny, Kristina Monakhova, Richard W. Shuai, and Laura Waller. 2022. Deep learning for fast spatially varying deconvolution. *Optica* 9, 1 (Jan 2022), 96–99. <https://doi.org/10.1364/OPTICA.442438>
- Ali Özgür Yöntem, Jinfeng Li, and Daping Chu. 2018. Imaging through a projection screen using bi-stable switchable diffusive photon sieves. *Optics express* 26, 8 (2018), 10162–10170.
- Tianjiao Zeng and Edmund Y. Lam. 2021. Robust Reconstruction with Deep Learning to Handle Model Mismatch in Lensless Imaging. *IEEE Transactions on Computational Imaging* 7 (2021), 1080–1092. <https://doi.org/10.1109/TCI.2021.3114542>
- Yucheng Zheng and M Salman Asif. 2021. Coded Illumination for Improved Lensless Imaging. *arXiv preprint arXiv:2111.12862* (2021).