



# Untrained networks for compressive lensless photography

KRISTINA MONAKHOVA,<sup>1,3,\*</sup>  VI TRAN,<sup>2,3</sup> GRACE KUO,<sup>1</sup>  AND LAURA WALLER<sup>1</sup> 

<sup>1</sup>*Department of Electrical Engineering & Computer Sciences, University of California, Berkeley, CA 94720, USA*

<sup>2</sup>*Orange Coast College, Costa Mesa, CA 92626, USA*

<sup>3</sup>*These authors contributed equally to this work*

\*[monakhova@berkeley.edu](mailto:monakhova@berkeley.edu)

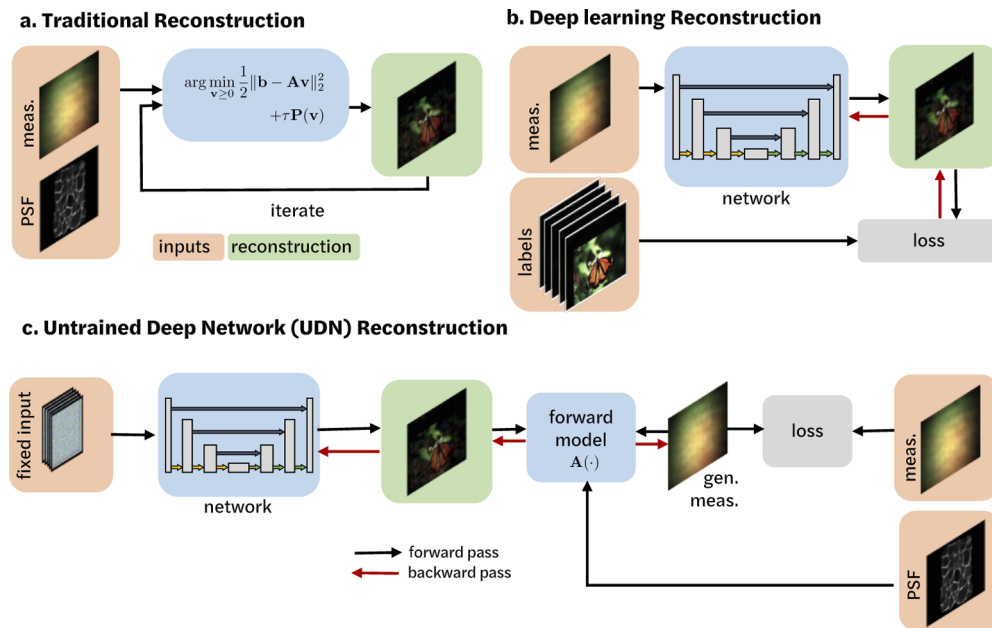
**Abstract:** Compressive lensless imagers enable novel applications in an extremely compact device, requiring only a phase or amplitude mask placed close to the sensor. They have been demonstrated for 2D and 3D microscopy, single-shot video, and single-shot hyperspectral imaging; in each case, a compressive-sensing-based inverse problem is solved in order to recover a 3D data-cube from a 2D measurement. Typically, this is accomplished using convex optimization and hand-picked priors. Alternatively, deep learning-based reconstruction methods offer the promise of better priors, but require many thousands of ground truth training pairs, which can be difficult or impossible to acquire. In this work, we propose an unsupervised approach based on untrained networks for compressive image recovery. Our approach does not require any labeled training data, but instead uses the measurement itself to update the network weights. We demonstrate our untrained approach on lensless compressive 2D imaging, single-shot high-speed video recovery using the camera's rolling shutter, and single-shot hyperspectral imaging. We provide simulation and experimental verification, showing that our method results in improved image quality over existing methods.

© 2021 Optical Society of America under the terms of the [OSA Open Access Publishing Agreement](#)

## 1. Introduction

Compressive imagers aim to recover more samples than they measure by leveraging compressive sensing, which guarantees signal recovery for underdetermined systems given certain assumptions about signal sparsity as well as incoherence in the measurement domain [1,2]. For optical imaging, compressive imagers have been demonstrated for many applications, including single-pixel and coded-aperture cameras [3–5]. These imagers have been shown to be effective for a number of compressive sensing tasks, such as single-shot lightfield imaging [6], 3D imaging [7], hyperspectral [8,9], and high-speed video imaging [10–13]. Recently, there has been a push to make compressive imagers more accessible by using inexpensive and compact hardware [14,15]. One such promising method includes lensless, mask-based imagers, which remove the lens of a traditional camera and instead place a phase or amplitude mask near the sensor to randomize the measurement, approximately achieving the measurement domain incoherence required for compressive sensing. These lensless mask-based cameras can be incredibly compact, with the mask only millimeters from the sensor, and assembly does not require precise alignment [16–18]. Mask-based lensless imagers have been demonstrated for compact single-shot 3D fluorescence microscopy [19–22], hyperspectral imaging [23], and high speed single-shot video [24]. These sorts of cameras are very promising for a number of imaging modalities; however, their performance depends on the fidelity of the reconstruction algorithm. Since the algorithm must solve an underdetermined inverse problem with assumptions on signal sparsity, the choice of algorithm and imaging priors used can significantly impact results.

Traditionally, images from compressive cameras are recovered by solving a convex optimization problem, minimizing both a least-squares loss based on the physics of the imaging system and a hand-chosen prior term, which enforces sparsity in some domain, Fig. 1(a). For successful image recovery by compressive sensing, there must be both incoherence in the sensing basis and sufficient sparsity in the sample domain. Lensless mask-based cameras utilize multiplexing optics to fulfill the incoherent sampling requirement, mapping each pixel in the scene to many sensor pixels in a pseudo-random way. The prior term enforces sparsity and ensures successful signal recovery. Over the years, a number of different hand-picked priors have been used for compressive imaging, such as sparsity in wavelets, total-variation (TV), and learned dictionaries. TV, which is based on gradient sparsity, has been particularly popular [25]. These methods are effective at solving imaging inverse problems, however they have recently been outperformed by deep learning-based methods, which incorporate non-linearity and network structures that may be better suited to image representation [26]. Recent work on plug and play (PnP) priors [27–30] offers some hope of combining inverse problems with state of the art denoisers (both deep and not, e.g. BM3D [31]); however, PnP still relies on using hand-picked denoisers or pre-trained networks that may not be well-suited for a given application.



**Fig. 1. Comparison of different reconstruction approaches.** (a) Traditional reconstructions solve a convex optimization problem with a data-fidelity-term based on the known forward model and a regularization term which acts as the image prior. (b) Deep learning methods require thousands or more labeled image pairs in order to train a neural network to approximate the reconstruction. (c) Our untrained deep network (UDN) uses a neural network as the image prior, but does not require any training data. The system forward model is used to calculate the loss between the estimated image and the measurement, which is then used to update the network weights.

Recently, several deep-learning based approaches have shown improved reconstruction quality for imaging inverse problems [26,32,33]; however, they rely on having a large dataset of experimental labeled ground truth and measurement pairs. In these methods, a deep neural network is used to approximate the imaging inverse problem, Fig. 1(b). The network is typically trained end-to-end and requires large datasets of labeled image and ground truth pairs. The

network can be agnostic to the imaging system physics, or the network can incorporate knowledge of the physics for faster convergence and reduced training data requirements [34,35]. For high-dimensional imaging applications, such as 3D and hyperspectral lensless imaging, obtaining ground truth datasets can be impractical, costly, or impossible. While there has been some work in training networks from synthetic data, these methods can suffer from model-mismatch if the synthetic data does not match the experimental system [36].

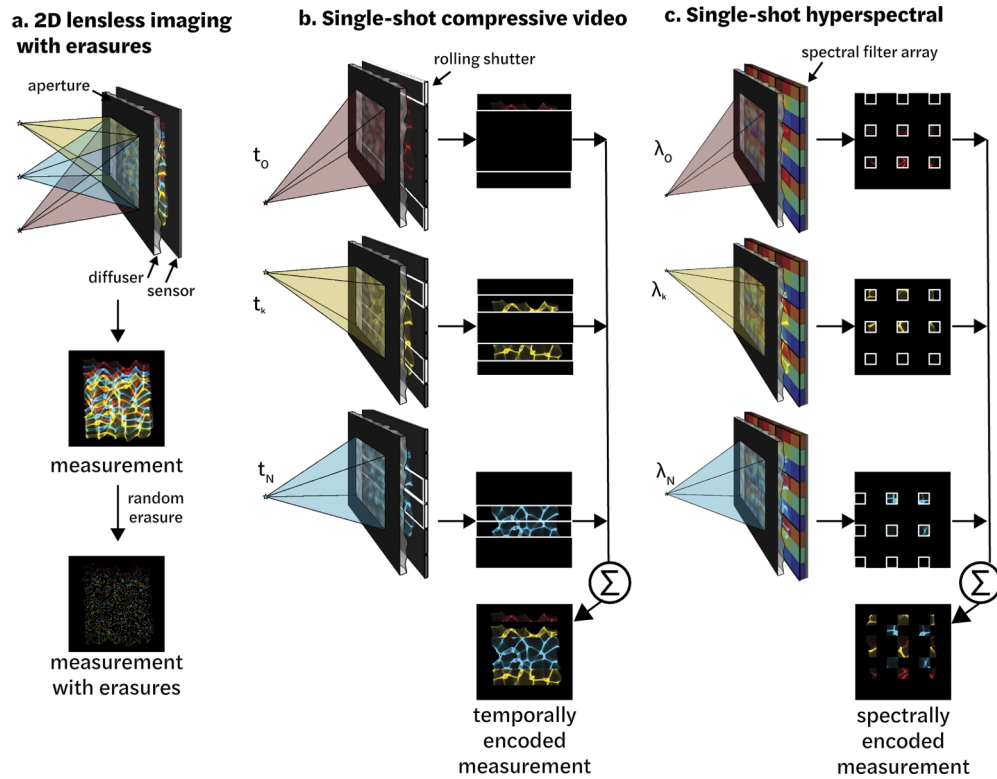
Recent work using unsupervised learning with untrained networks is especially promising for a number of imaging applications - leveraging the structure of neural networks without needing any training data. Untrained networks, such as deep image prior [37] and deep decoder [38], have shown that the structure of neural networks can be effective at serving as a prior on image statistics without any training. An untrained deep network with randomly initialized weights is used as an image generator, outputting the recovered image. The network weights are updated through a loss function comparing the generated image with the input data, for example, a noisy image. This method and several related papers have been shown to be particularly effective for simulated image denoising, deblurring, and super-resolution [39,40]. For many computational imaging problems, the measurements do not resemble the reconstructed image; instead, the scene is related to the measurement through a forward model that describes the physics of the image formation problem. For this class of problems, untrained networks can be paired with differentiable imaging forward models in which the network weights are updated based on a loss between the measurement and the generated measurement from the network output passed through the imaging forward model, Fig. 1(c). This has been shown to be effective on several imaging modalities, such as phase imaging [41], MRI [42–44], diffraction tomography [45], and several small-scale compressive sensing problems [46]. Here, we extend this framework to compressive lensless photography.

We propose a reconstruction method for compressive lensless imaging based on untrained networks, which we call untrained deep network (UDN) reconstructions. Our approach uses a deep network for the image prior, but requires no training data. We present a general differentiable imaging model that could be used for multiple types of compressive lensless imaging, and test it on three different cases: 2D lensless imaging with erasures, single-shot video, and single-shot hyperspectral imaging.

To the best of our knowledge, untrained networks have not been used for compressive optical imaging before, thus providing both a stress test of untrained deep networks on several challenging underdetermined experimental systems, as well as bringing improved image quality to compressive lensless imaging. We provide simulation and experimental results, showing improved performance over existing methods that do not utilize training data. Our results indicate that untrained networks can serve as effective image priors for these systems, providing better reconstructions in cases where it is not possible to obtain labeled ground truth data.

## 2. Imaging models

Our system architecture is based on the lensless mask-based camera, DiffuserCam [17,47], which consists of a thin diffuser (a smooth pseudorandom phase optic) which is placed a few millimeters in front of the sensor, Fig. 2(a). Light from each point in the scene refracts through the diffuser to create a high-contrast pseudorandom caustic pattern on the sensor plane. Since each point in the world maps to many pixels on the sensor, we need only a subset of the sensor pixels to reconstruct the full 2D image, given sufficient sparsity. This can be useful, for example, to compensate for dead sensor pixels, which act like an erasure pattern; or, we can capture the full 2D sensor image and use compressed sensing to recover higher-dimensional information. The higher-dimensional data (e.g. time, wavelength) must be physically encoded into different subsets of the sensor pixels; then compressed sensing is applied to recover a 3D datacube.



**Fig. 2. Imaging systems:** Our lensless 2D camera consists of a diffuser placed a small distance away from the sensor. Each point in the world maps to a high-contrast caustic pattern (PSF). (a) Since the PSF covers a large area of the sensor, we can erase a random subset of the pixels and still recover the full image. (b) The camera's rolling shutter can be used to encode temporal information into the measurement. As points flick on/off, the PSFs are filtered by the sensor's rolling shutter function, which reads one row at a time (or two in the case of dual shutter, as shown here). The measurement is a sum of the rows captured at different time points, each of which contains information from the entire 2D scene. Hence, temporal information is encoded vertically across the sensor, with earlier time points at the top and bottom of the image and later time points towards the center. (c) Single-shot hyperspectral imaging is achieved with a spectral filter array in front of the sensor, which maps each band of wavelengths to a subset of the sensor pixels.



We demonstrate our UDN on two examples of three-dimensional data recovery from a single-shot 2D measurement: (1) high-speed video [24] and (2) hyperspectral imaging [23]. The high-speed video example takes advantage of the sensor's built-in rolling shutter, which exposes different rows of pixels at each time point, Fig. 2(b), effectively acting like an erasure pattern at each time point. From this data, one can use compressed sensing to recover a separate 2D image for each time point; the result is a full-resolution video at the framerate of the rolling shutter (the line scan rate). For the hyperspectral imaging example, the DiffuserCam device has an added spectral filter array in front of the sensor, with 64 different color channels, Fig. 2(b). Each of the 64 wavelengths thus map to a different subset of the sensor pixels, similarly acting like an erasure pattern for each wavelength. From this data, one can use compressed sensing to recover a hyperspectral datacube (2D spatial + 1D spectral) from a single 2D measurement. In both examples, we choose to reconstruct the full datacube simultaneously, rather than a sequential set of 2D reconstructions, since this allows us to add priors along the time/wavelength dimension.

We now go into detail on the modeling, starting with the 2D image formation model of DiffuserCam, and building to our three example scenarios: 2D imaging with erasures, high-speed video, and hyperspectral imaging. These formulations will be used to create the differential forward model within our network architecture.

### 2.1. 2D lensless imaging

As described in [17,47], we model our lensless camera as shift-invariant: we assume the pattern on the sensor from each point in the world is a translated version of the on-axis pattern, called the point spread function (PSF). We can therefore model the sensor response  $\mathbf{b}[x, y]$  as a convolution of the scene  $\mathbf{v}[x, y]$  with an on-axis PSF  $\mathbf{h}[x, y]$ :

$$\mathbf{b}[x, y] = \text{crop}\left(\mathbf{v}[x, y] * \mathbf{h}[x, y]\right) \quad (1)$$

$$= \mathbf{A}_{2D} \mathbf{v}. \quad (2)$$

where  $[x, y]$  represents the discrete image coordinates,  $*$  represents a discrete 2D linear convolution and the crop function accounts for the finite sensor size. For 2D imaging, we assume that objects are placed beyond the hyperfocal distance of the imager so that the PSF does not vary with depth. In addition, we assume that there is no wavelength variation in the PSF as demonstrated in [23].

### 2.2. 2D imaging with erasures

For the case of 2D imaging with erasures, we multiply the result of Eq. (1) with a binary mask, denoted  $\mathbf{M}[x, y]$ , that zeros out a subset of the sensor pixels to model the erasure pattern:

$$\mathbf{b}[x, y] = \mathbf{M}[x, y] \cdot \text{crop}\left(\mathbf{v}[x, y] * \mathbf{h}[x, y]\right) \quad (3)$$

$$= \mathbf{A}_{2D \text{ erasures}} \mathbf{v}. \quad (4)$$

As we show in Sec. 5., compressed sensing enables recovery of a full image, even in the presence of a significant fraction of dead or missing pixels. Although commercial sensors are screened to minimize dead pixels, this scenario is useful for exploring the limits of compressive sensing in lensless cameras since we can synthetically increase the percentage of erasures to increase the ratio of data reconstructed to data measured.

### 2.3. Higher dimensions: video and hyperspectral

Rather than sampling only a subset of pixels in order to reconstruct a 2D image, compressed sensing can instead be used to reconstruct 3D datasets from fully-sampled 2D images. If the higher-dimensional data is physically encoded into different subsets of the pixels, one can use

Eq. (3) to recover a collection of 2D images from a single acquisition. As described previously, our single-shot video example uses the rolling shutter to encode temporal information into different pixels, and our hyperspectral example uses a spectral filter array to encode wavelength information. We denote this extra dimension (time/wavelength) generically as the  $k$ -dimension. Sequential recovery using Eq. (3) prevents incorporating priors along the  $k$ -dimension, so we use the following model that depends on the full datacube:

$$\mathbf{b} = \sum_{k=0}^{N_k} \mathbf{M}_k[x, y] \cdot \text{crop}\left(\mathbf{v}[x, y, k] * \mathbf{h}[x, y]\right) \quad (5)$$

$$= \mathbf{A}_{k-D} \mathbf{v}. \quad (6)$$

Here,  $N_k$  is the number of discrete points along the  $k$ -dimension, and  $\mathbf{M}_k[x, y]$  is a masking function, which depends on  $k$ , and selects the sensor pixels corresponding to each video frame/wavelength. The convolution,  $*$  is only over the two spatial dimensions.

For the high-speed video case  $\mathbf{M}_k[x, y]$ , referred to as the shutter function, is based on the rolling shutter. Sensors typically have either a single shutter or a dual shutter, which we use in this work. For a single shutter, the sensor reads the pixels one horizontal line at a time, moving from the top to the bottom of the sensor; for a dual-shutter, the sensor reads pixels from two horizontal lines that move from the top and bottom of the sensor towards the middle, Fig. 2(b).

For the hyperspectral case  $\mathbf{M}_k[x, y]$ , referred to as the filter function, is determined by the spectral filter array. Each filter pixel acts as narrow-band spectral filter which integrates light within a certain wavelength range and blocks out light at other wavelengths. We approximate this as a finite sum across spectral bands with a non-binary filter function.

### 3. Inverse problem

Given our sensor measurement  $\mathbf{b}$ , our goal is to recover the scene  $\mathbf{v}$ . For the 2D erasures scenario  $\mathbf{v}$  is a 2D image, whereas for single-shot video  $\mathbf{v}$  is a video consisting of two spatial and one temporal dimension, and for single-shot hyperspectral  $\mathbf{v}$  is a hyperspectral volume consisting of two spatial and one spectral dimension. In all cases, the problem is underdetermined, since we aim to recover more than we measure. First, we describe the traditional reconstruction methods based on convex optimization which will serve as our baseline comparison, and then we describe our untrained network reconstruction method.

#### 3.1. Traditional inverse problem

Due to incoherence of the measurement system that comes from our use of a multiplexing phase optic (diffuser), compressive sensing can be utilized to recover the image. By compressive sensing theory, we can formulate our inverse problem as follows:

$$\hat{\mathbf{v}} = \arg \min_{\mathbf{v} \geq 0} \frac{1}{2} \|\mathbf{b} - \mathbf{A}\mathbf{v}\|_2^2 + \tau R(\mathbf{v}), \quad (7)$$

where  $R(\mathbf{v})$  is a prior on the scene, and is often of the form  $\|\mathbf{D}\mathbf{v}\|_1$ , where  $\mathbf{D}$  is a sparsifying transform.  $\mathbf{A}$  is our forward model, which can be of the form  $\mathbf{A}_{2D}$ ,  $\mathbf{A}_{2D \text{ erasures}}$ , or  $\mathbf{A}_{k-D}$ .

In practice, 2D or 3D TV priors work well for a variety of scenes, and are implemented by defining the regularizer term as  $R(\mathbf{v}) = \|\nabla_{xyk} \mathbf{v}\|_1$ , where  $\nabla_{xyk} = [\nabla_x \nabla_y \nabla_k]^T$  is the matrix of forward finite differences in the  $x$ ,  $y$ , and  $k$  directions. Convex optimization approaches such as fast iterative shrinkage-thresholding algorithm (FISTA) [48] or alternating direction method of multipliers (ADMM) [49] can be used to efficiently solve this problem. In each case, the prior is hand-tuned for the given application by varying the amount of regularization through the tuning parameter  $\tau$ , which trades off data-fidelity and regularization. As a baseline comparison, we use

FISTA with 2DTV for our 2D imaging with erasures problem and weighted anisotropic 3DTV for our single-shot video and single-shot hyperspectral imaging [50]. We include an additional comparison using PnP-BM3D, PnP-BM4D [51], and a pretrained PnP denoiser network [29] in Supplement 1.

### 3.2. Untrained network

In this work, we propose to instead use an untrained network for the image reconstruction. It has been shown that neural networks are good at representing and generating natural images; thus, minimizing an image generator network instead of directly solving for the image could be a good way to effectively regularize with a custom deep prior.

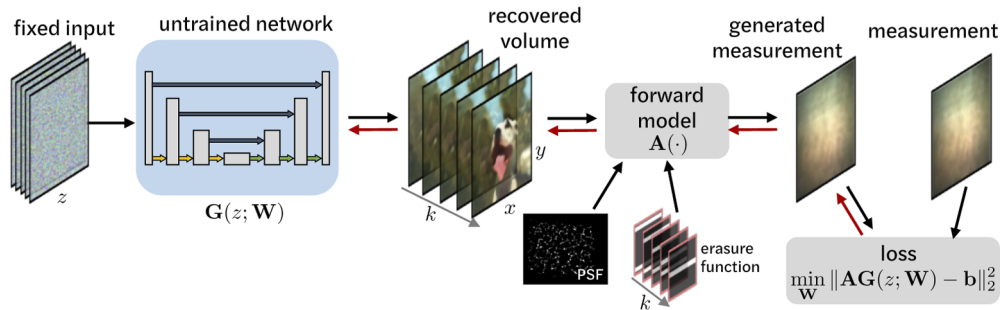
Rather than solving for the image  $\mathbf{v}$  directly as before, we solve for the image indirectly as the output of the generator network  $G(z; \mathbf{W})$ . This output image  $\mathbf{v}_{gen}$  is then passed through our imaging forward model  $\mathbf{A}$  and compared against the measured image  $\mathbf{b}$ . The network weights are updated based on the difference between the generated measurement and the true measurement, which corresponds to solving the following problem:

$$\mathbf{W}^* = \arg \min_{\mathbf{W}} \frac{1}{2} \|\mathbf{b} - \mathbf{A}\mathbf{G}(z; \mathbf{W})\|_2^2 \quad (8)$$

$$= \arg \min_{\mathbf{W}} \frac{1}{2} \|\mathbf{b} - \mathbf{A}\mathbf{v}_{gen}\|_2^2, \quad (9)$$

where our network  $G(z; \mathbf{W})$  has a fixed input  $z$  and randomly initialized weights  $\mathbf{W}$ . The network output  $\mathbf{v}_{gen}$  is the reconstructed image.  $z$  can be thought of as a latent code to our generator network and is randomly initialized and held constant. We update the weights of the network via backpropagation in order to minimize the loss between the actual measurement and the generated measurement (Fig. 3). This process must be repeated for each image reconstruction, since there are no ‘training’ and ‘testing’ phases as there are for deep learning methods with labeled data.

#### Untrained Deep Network (UDN) Reconstruction



**Fig. 3. Overview of our Untrained Deep Network (UDN) reconstruction pipeline.** An untrained network with randomly initialized weights takes in a fixed, random input vector. The network outputs a sequence of images along the  $k$ -axis (for video imaging, one image for each time point), and we pass this output volume into our known imaging forward model (defined by the PSF calibration image and known erasure function) to generate a single simulated measurement. We compare the generated measurement with our captured measurement and use the difference between the two to update the untrained network parameters.

As in [37], we utilize an encoder-decoder architecture with skip connections and keep the input to the network fixed. See Supplement 1 for details on our network architecture and hyperparameters for each experiment.

#### 4. Implementation details

For the 2D erasures case, we utilize an existing 2D lensless imaging dataset which consists of pairs of 2D lensless camera measurements and corresponding lensed camera ground truth images [34]. This dataset is advantageous because it includes experimental lensless measurements with real sensor noise and other non-idealities, but also provides labeled ground truth data from an aligned lensed camera. For our simulation results, we utilize the ground truth images from the dataset to generate simulated measurements using the forward model in Eq. (3); for our experimental results, we directly use the lensless camera measurements. For both, we synthetically add erasures to the measurement by point-wise multiplying it with an erasure function with 0%, 50%, 90%, 95%, and 99% erasures, picking a random subset of indices to erase.

For single-shot video and single-shot hyperspectral imaging, there does not exist experimental data with associated ground truth. For our experimental data analysis, we utilize experimental raw sensor measurements from [24] and [23]. This includes measurements taken from a PCO Edge 5.5 sCMOS camera with a dual rolling shutter and a homemade random diffuser (for single-shot video), and measurements from a hyperspectral DiffuserCam with a spectral filter array and Luminit diffuser [52] (for hyperspectral imaging). For our simulations, we use the PSF and shutter/filter functions from these systems to simulate our sensor measurements.

We implement our network and differentiable forward model with mean square error (MSE) loss in Pytorch and run our reconstructions on a Titan X GPU with the ADAM optimizer throughout training. We perform early stopping to obtain the best reconstructions, as described in [37], with reconstructions ranging from 1,000 - 100,000 iterations, which generally takes several hours. This process must be completed for every new image; unlike with deep learning methods that use training data, there is no distinction between training and testing and instead the network parameters must be re-optimized for every reconstruction. For reproducibility, training code will be available here upon publication of the paper: <https://github.com/Waller-Lab/UDN/>.

#### 5. Results

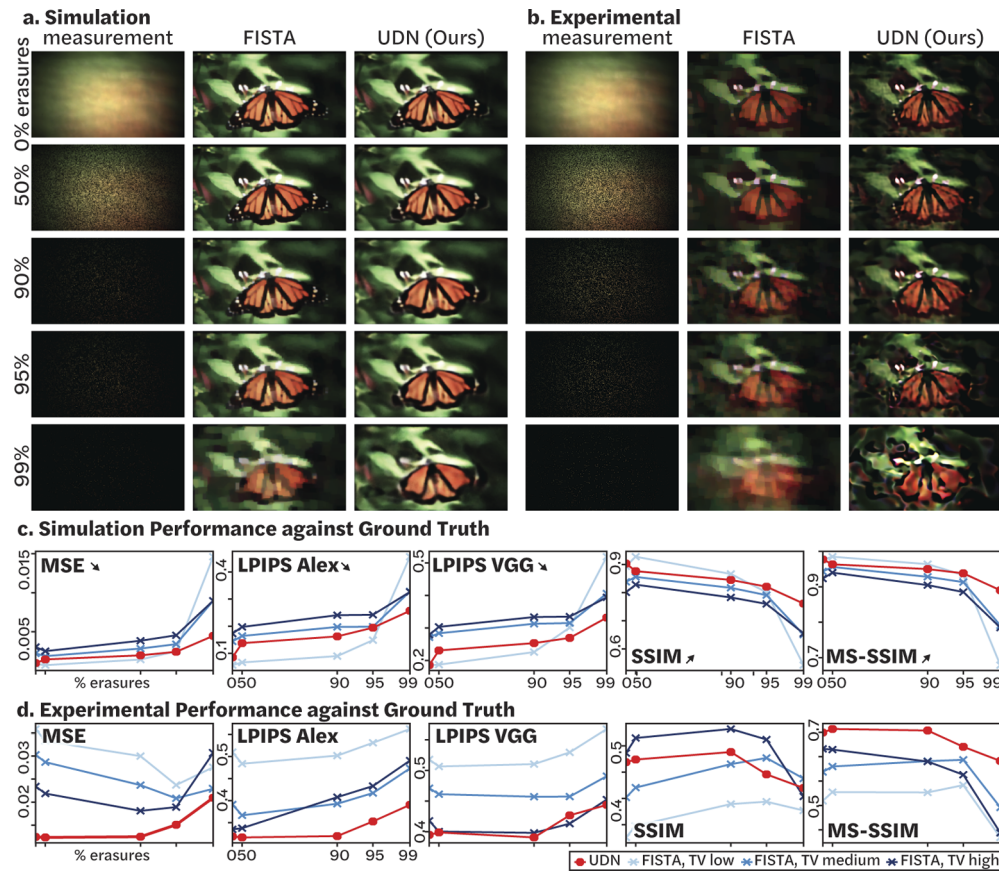
We compare the results of our untrained reconstruction method against the standard FISTA reconstruction with TV regularization. For all three cases, we present both simulation results as well as experimental validation. For the FISTA reconstructions, we reconstructed images with three different amounts of TV regularization, since more regularization leads to improved image quality and lower mean squared error, but tends to blur high-frequency information. Less regularization reveals high-frequency information, but at the price of increased reconstruction artifacts. When ground truth is available (in all simulations and in experiment for 2D erasures), we compare our reconstructed images to the ground truth via a variety of quality metrics: mean squared error (MSE), learned perceptual similarity metric (LPIPS) based on AlexNet and VGG [53], the structural similarity index measure (SSIM), and the multi-scale structural similarity measure (MS-SSIM) [54]. MSE is a standard image metric, but tends to favor low-frequency information. SSIM and MS-SSIM are both perception-based metrics, with MS-SSIM using multiple image scales and achieving better perceptual performance than SSIM. Both LPIPS metrics are learned perceptual similarity metrics based on deep networks, with LPIPS VGG being closer to a traditional perceptual similarity metric. Each method has its strengths and weaknesses, and therefore we provide comparisons using each of the metrics when possible (MS-SSIM requires a certain minimum image size which precludes us from using it for the single-shot video and single-shot hyperspectral cases, and LPIPS only works on RGB color images, so we cannot use it for single-shot hyperspectral). For MSE, LPIPS Alex, and LPIPS VGG, a lower score is better, whereas for SSIM and MS-SSIM, a higher score is better.



## 5.1. 2D compressive imaging

### 5.1.1. Simulation

First, we simulate a noise-free 2D measurement with increasing numbers of randomly-distributed pixel erasures (0%, 50%, 90%, 95%, and 99%), then compare reconstruction results using both FISTA and our UDN (Fig. 4). FISTA uses high, medium, and low amounts of TV corresponding to  $\tau = 1e-4, 5.5e-4, 1e-5$ , respectively.



**Fig. 4. 2D Compressive imaging with erasures: simulation and experimental results.**

Reconstruction results for increasing numbers of pixel erasures, showing the measurement along with the best fast iterative shrinkage-thresholding algorithm (FISTA) reconstruction (based on the LPIPS Alex metric) and the UDN reconstruction for (a) simulated measurements and (b) experimental measurements with synthetically added erasures. (c) and (d) compare our performance against the ground truth for several quality metrics (arrows indicate which direction is better), showing that in simulation UDN outperforms FISTA only at 99% erasures, whereas in experiments with real noise and non-idealities, UDN outperforms FISTA on the perceptual image metrics, LPIPS Alex and MS-SSIM, as well as on MSE.

We compare our performance against the ground truth images on our five image quality metrics for an increasing percentage of erasures in Fig. 4(b) and show the corresponding images in Fig. 4(a) for our method vs. the best FISTA result (based on the LPIPS-Alex metric). From this, we can see that FISTA and our method perform similarly well for 0%-95% erasures, but our method has improved performance for 99% erasures. This shows that TV regularization is a sufficient prior for smaller numbers of erasures in the absence of noise or model non-idealities,



but as the problem becomes severely underdetermined, our UDN provides a better image prior than TV and leads to improved image quality on all of our performance metrics.

### 5.1.2. Experimental

Using experimentally captured images from [34], we synthetically add random pixel erasures to the measured data and compare our reconstructions against the ground truth lensed images using the five image metrics described above. As a baseline, we compare against FISTA with high, medium, and low amounts of TV ( $\tau = 1e-1, 1e-2, 1e-3$ ), shown in Fig. 4(d). A visual comparison with the best FISTA reconstruction (based on the LPIPS-Alex metric) is shown in Fig. 4(c). Unlike in the noise-free simulation, our method consistently performs better than FISTA for all amounts of erasures on the MSE, LPIPS Alex, and MS-SSIM metric, resulting in a clearer and sharper reconstruction. Thus, our method gives better performance improvements over FISTA for real-world datasets, likely because the UDN provides a better image prior and outperforms TV in the presence of real measurement noise and imaging non-idealities.

TV is a very commonly used prior that is computationally efficient, but other priors can be incorporated using the PnP framework. See [Supplement 1](#) for additional comparisons against PnP-BM3D and a PnP pretrained denoiser network, demonstrating that the UDN outperforms these methods as well.

## 5.2. *Single-shot video*

### 5.2.1. Simulation

Using an experimentally captured PSF and shutter function from [24], we simulate a compressive video measurement, shown in Fig. 5 top left, using a 38-frame video. Fig. 5 shows the results of our reconstruction compared to the FISTA result for several frames of the video (with  $\tau = 1e-2, 1e-3, 5e-4$  for FISTA). It is evident that our method generates a more visually appealing result with fewer artifacts, while preserving high-resolution features throughout the frames. This is quantified in Fig. 5(bottom), showing that UDN has a significantly better performance on all metrics. For each of the FISTA results, the values are worse at the beginning and end of the video sequence, which is consistent with [24]. Our method similarly has worse performance for the first and last frames, but the difference is not as pronounced, resulting in more uniform image quality throughout the recovered video.

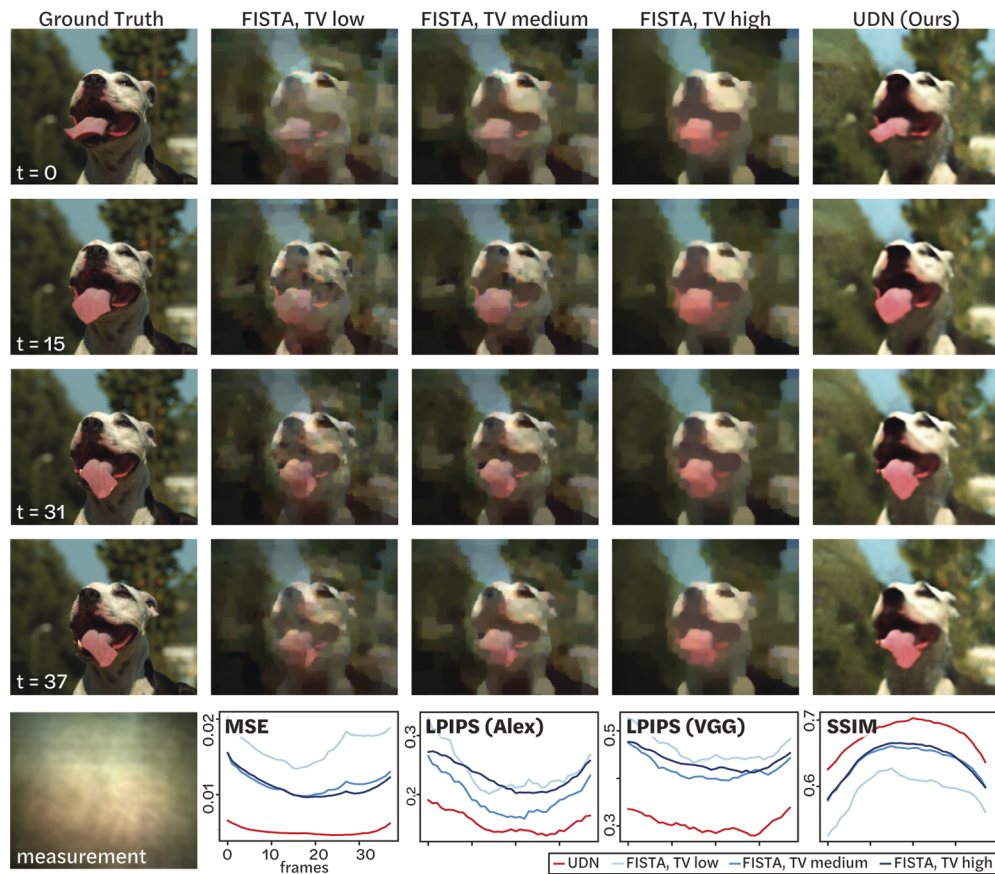
### 5.2.2. Experimental

Next, we recover 72-frame videos from a single experimental measurement, Fig. 6. When compared to the FISTA reconstructions (with  $\tau = 1e-2, 1e-3, 1e-4$ ), we can see that our method has more uniform image quality throughout. The foam dart has significant artifacts in the FISTA reconstruction for the first and last scenes of the video, even disappearing for low TV. Our method appears to capture the dart motion throughout the video frames and has fewer noticeable artifacts than the FISTA reconstruction. There is no ground truth available for this data, but our method seems to produce a more realistic and visually appealing reconstruction than FISTA in terms of the dart motion. FISTA with smaller amounts of TV is able to better recover the text on the book in the scene, but at the expense of reduced image quality and more artifacts throughout the video. Our method is not quite able to recover the book text, but has more uniform image quality and better captures the motion in the scene.


## 5.3. *Single-shot hyperspectral*

### 5.3.1. Simulation

Finally, we recover a hyperspectral volume with 64 wavelengths from a single simulated measurement. We simulate the measurement using an experimentally captured PSF and filter



**Fig. 5. Simulation results on single-shot video.** We recover a 38-frame video from a single measurement (bottom left). We show four sample reconstructed video frames in the top four rows, and plot the quality metrics for all frames below. Here we see that our reconstruction has sharper features across frames, enabling superior recovery especially for the first and last frames, where FISTA has more pronounced reconstruction artifacts. We achieve better MSE, LPIPS, and SSIM scores across frames (bottom right). See [Visualization 1](#) for the full video.

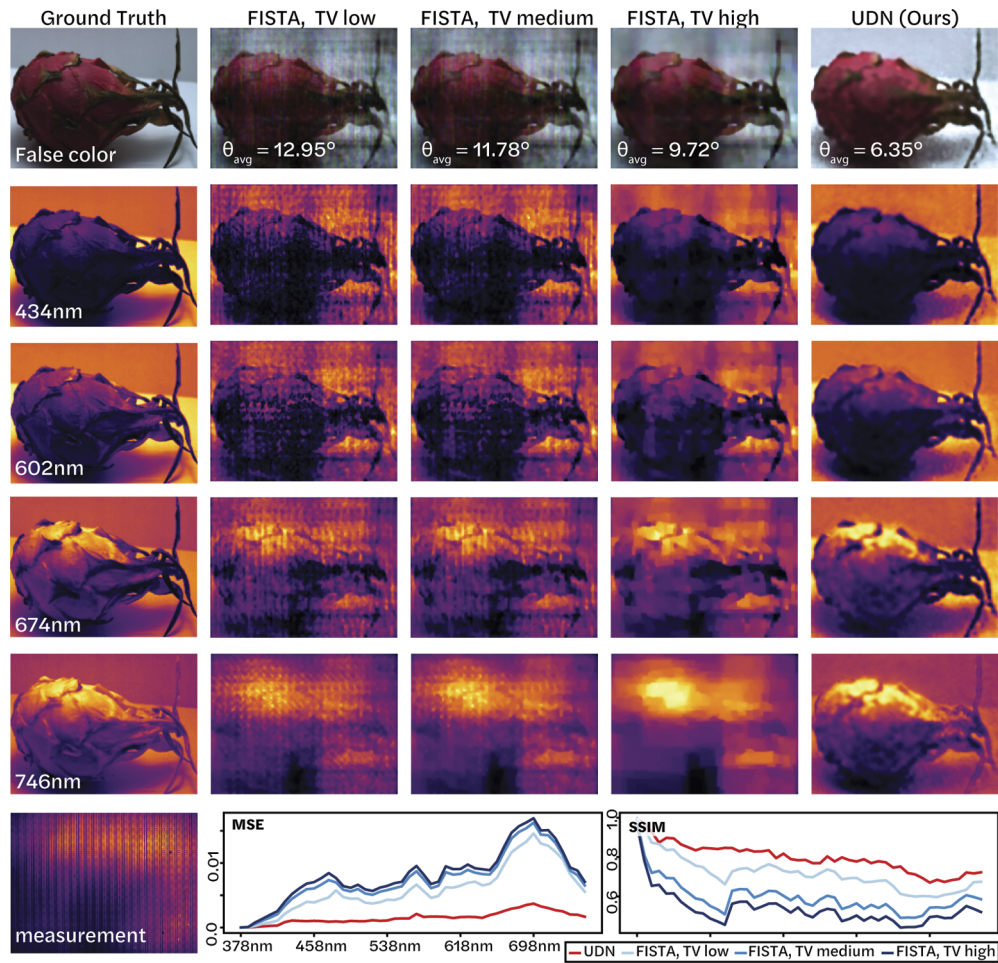


**Fig. 6. Experimental results on single-shot video.** We recover 72 frames from a single measurement (top left). The rows show four sample reconstructed frames from our 72-frame reconstruction, with both our UDN method and FISTA with three different amounts of TV. Here we see that our reconstruction has sharper features across frames and better captures motion within the video. See [Visualization 2](#) for the full video and [Visualization 3](#) for a second experimental example.

function from [23], along with a ground truth hyperspectral image from [55] with 64 spectral channels. Figure 7 shows the results of our reconstruction compared to FISTA with high, medium, and low amounts of TV ( $\tau = 3e-7$ ,  $6e-7$ , and  $3e-6$ ). We can see that UDN preserves more features across wavelengths and has fewer artifacts than FISTA. We compare the MSE and SSIM across the methods, Fig. 7(bottom). We can see that UDN has significantly better MSE and SSIM values than FISTA across wavelengths. In addition, we report the average cosine distance between the spectral profiles in the reconstruction and ground truth images (Fig. 7). The cosine distance is especially important for hyperspectral imaging due to its role in hyperspectral classification. UDN provides a lower average cosine distance than FISTA, indicating that UDN achieves a better recovery of the spectral profiles. We note that FISTA with high TV achieves a better cosine distance than FISTA with low TV, but at the price of worse spatial resolution. Our UDN method achieves both better spatial quality (MSE and SSIM) as well as better spectral performance (cosine distance) than FISTA.

### 5.3.2. Experimental


We test our performance on an experimentally captured measurement [23] of a Thorlabs plush dog illuminated by a broadband lightsource, recovering 32 spectral channels from a single measurement (downsampled  $2\times$  spectrally). We compare against FISTA with low and high TV ( $\tau = 3e-4$  and  $5e-5$ ) along with the best reconstruction from [23] (FISTA TV + low rank),



**Fig. 7. Simulation results on hyperspectral data.** We display a false-color image of the recovered 64-channel hyperspectral volume (top row), along with four selected spectral slices. The quality metrics are plotted for each wavelength at bottom right. Here we see that our reconstruction has sharper features and fewer artifacts than FISTA, and achieves a better MSE and SSIM score across all wavelengths. In addition, UDN achieves better cosine similarity ( $\theta_{\text{avg}}$ ) to the ground truth spectra than FISTA. See [Visualization 4](#) for the full hyperspectral volume.



down-sized to match our reconstruction size. While ground truth images do not exist for this data, UDN appears to provide more consistent image quality and retains more of the details across wavelengths than FISTA, Fig. 8.



**Fig. 8. Experimental results on hyperspectral data.** We show a false-color image of the recovered 32-channel hyperspectral volume (top row), along with four spectral slices. Here we see that our reconstruction has sharper features and fewer artifacts across wavelengths. See [Visualization 5](#) for the full hyperspectral volume.

## 6. Discussion

Untrained networks offer a number of distinct advantages for compressive lensless imaging. First, they do not require any training data, as opposed to deep learning-based methods. Training data is especially hard or impossible to acquire for higher-dimensional imaging, such as for high-speed video, hyperspectral, or 3D imaging, so this feature is particularly useful. Untrained networks can serve as a better prior for certain high-dimensional imaging problems off-the-shelf, potentially enabling better reconstructions for a number of compressive imaging modalities.

Currently, the main limitations of untrained networks are: 1) memory-constraints and 2) speed. First, many of our reconstructions are GPU memory limited. To take advantage of accelerated GPU processing, the entire untrained network must fit in memory on the GPU, limiting the size of reconstructions that we can process, and we find that we can process much larger images and volumes with FISTA than we can with UDN. In this work, our single-shot video and single-shot hyperspectral measurements are down-sized between 2-16 $\times$  from the original size in order to fit in the GPU, limiting our resolution. Looking forward, larger GPU sizes or clever computational



techniques to better utilize GPUs for memory-limited problems could improve this and enable the reconstruction of larger images and volumes. Next, the speed of the untrained network reconstructions is generally an order-of-magnitude slower than standard FISTA reconstructions. Thus, our method is best suited for applications where a real-time reconstruction is not needed, since typical untrained reconstructions take between 1-5 hours. As machine learning speeds and processors improve, we expect these factors to be less limiting.

Given the benefits and limitations of untrained networks, we envision this reconstruction method to be useful for a certain subset of problem where it is difficult or impossible to obtain ground truth data for training, but where there are little or no time constraints on the reconstruction. Given the fact that no training data is needed, this method can be applied to many different imaging problems without needing to collect new datasets. UDNs are especially promising for imaging dense, natural scenes where TV can be a poor prior.

One interesting open problem for untrained reconstructions is the network choice. In our experiments, we utilized a convolutional network with an encoder-decoder structure and found that this worked well; however, our network tended to blur out high frequency features and there may be other architectures that could potentially provide better priors. For instance, deep decoder [38] has been demonstrated for similar tasks, but we found that for our application it required more iterations to converge and was outperformed by an encoder-decoder structure, demonstrating that the choice of network architecture is important and application-specific. Although our network worked well for photographic scenes, other network architectures may be more appropriate for other types of images, such as fluorescent biological targets which may have very different statistics than photographic scenes.

## 7. Conclusion

We have demonstrated that untrained networks can improve the image quality for lensless compressive imaging systems. We tested untrained networks on 2D lensless imaging with varying amounts of erasures, and we demonstrated their effectiveness on single-shot compressive video and single-shot hyperspectral imaging, in which we recover a full 72-frame video or 32 to 64 spectral slices, respectively, from a single measurement. In each case, we showed both in simulation and experiment that untrained networks can have better image quality than compressive-sensing-based minimization using total-variation regularization, demonstrating that non-linear networks can be a better prior than TV for dense, natural scenes. We believe that untrained networks are especially promising for situations in which training data is difficult or impossible to obtain, providing a better imaging prior for underdetermined reconstructions.

**Funding.** Chan Zuckerberg Initiative; National Science Foundation (DGE 1752814, DMR 1548924); Gordon and Betty Moore Foundation (GBMF4562).

**Acknowledgments.** Vi Tran acknowledges support from the Transfer-to-Excellence REU program, funded by the Hopper Dean Foundation and hosted by the Center for Energy Efficient Electronics Science (NSF Award 0939514).

**Disclosures.** The authors declare no conflicts of interest

**Data availability.** All data will be available [38] upon publication of the paper.

**Supplemental document.** See [Supplement 1](#) for supporting content.

## References

1. E. J. Candès, "Compressive sampling," in *Proceedings of the international congress of mathematicians*, vol. 3 (European Mathematical Society, 2006), pp. 1433–1452.
2. E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.* **25**(2), 21–30 (2008).
3. M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE Signal Process. Mag.* **25**(2), 83–91 (2008).

4. M. B. Wakin, J. N. Laska, M. F. Duarte, D. Baron, S. Sarvotham, D. Takhar, K. F. Kelly, and R. G. Baraniuk, "An architecture for compressive imaging," in *2006 international conference on image processing*, (IEEE, 2006), pp. 1273–1276.
5. G. Huang, H. Jiang, K. Matthews, and P. Wilford, "Lensless imaging by compressive sensing," in *2013 IEEE International Conference on Image Processing*, (IEEE, 2013), pp. 2101–2105.
6. K. Marwah, G. Wetzstein, Y. Bando, and R. Raskar, "Compressive light field photography using overcomplete dictionaries and optimized projections," *ACM Trans. Graph.* **32**(4), 1–12 (2013).
7. A. Levin, R. Fergus, F. Durand, and W. T. Freeman, "Image and depth from a conventional camera with a coded aperture," *ACM Trans. Graph.* **26**(3), 70 (2007).
8. A. Wagadarikar, R. John, R. Willett, and D. Brady, "Single disperser design for coded aperture snapshot spectral imaging," *Appl. Opt.* **47**(10), B44–B51 (2008).
9. M. E. Gehm, R. John, D. J. Brady, R. M. Willett, and T. J. Schulz, "Single-shot compressive spectral imaging with a dual-disperser architecture," *Opt. Express* **15**(21), 14013–14027 (2007).
10. Y. Hitomi, J. Gu, M. Gupta, T. Mitsunaga, and S. K. Nayar, "Video from a single coded exposure photograph using a learned over-complete dictionary," in *2011 International Conference on Computer Vision*, (IEEE, 2011), pp. 287–294.
11. D. Reddy, A. Veeraraghavan, and R. Chellappa, "P2c2: Programmable pixel compressive camera for high speed imaging," in *CVPR 2011*, (IEEE, 2011), pp. 329–336.
12. P. Llull, X. Liao, X. Yuan, J. Yang, D. Kittle, L. Carin, G. Sapiro, and D. J. Brady, "Coded aperture compressive temporal imaging," *Opt. Express* **21**(9), 10526–10545 (2013).
13. L. Gao, J. Liang, C. Li, and L. V. Wang, "Single-shot compressed ultrafast photography at one hundred billion frames per second," *Nature* **516**(7529), 74–77 (2014).
14. D. S. Jeon, S.-H. Baek, S. Yi, Q. Fu, X. Dun, W. Heidrich, and M. H. Kim, "Compact snapshot hyperspectral imaging with diffracted rotation," *ACM Trans. Graph.* **38**(4), 1–13 (2019).
15. A. Liutkus, D. Martina, S. Popoff, G. Chardon, O. Katz, G. Lerosey, S. Gigan, L. Daudet, and I. Carron, "Imaging with nature: Compressive imaging using a multiply scattering medium," *Sci. Rep.* **4**(1), 5552 (2015).
16. M. S. Asif, A. Ayremlou, A. Sankaranarayanan, A. Veeraraghavan, and R. G. Baraniuk, "Flatcam: Thin, lensless cameras using coded aperture and computation," *IEEE Trans. Comput. Imaging* **3**(3), 384–397 (2017).
17. N. Antipa, G. Kuo, R. Heckel, B. Mildenhall, E. Bostan, R. Ng, and L. Waller, "DiffuserCam: lensless single-exposure 3D imaging," *Optica* **5**(1), 1–9 (2018).
18. T. Shimano, Y. Nakamura, K. Tajima, M. Sao, and T. Hoshizawa, "Lensless light-field imaging with fresnel zone aperture: quasi-coherent coding," *Appl. Opt.* **57**(11), 2841–2850 (2018).
19. J. K. Adams, V. Boominathan, B. W. Avants, D. G. Vercosa, F. Ye, R. G. Baraniuk, J. T. Robinson, and A. Veeraraghavan, "Single-frame 3D fluorescence microscopy with ultraminiature lensless FlatScope," *Sci. Adv.* **3**(12), e1701548 (2017).
20. G. Kuo, F. L. Liu, I. Grossrubatscher, R. Ng, and L. Waller, "On-chip fluorescence microscopy with a random microlens diffuser," *Opt. Express* **28**(6), 8384–8399 (2020).
21. K. Yanny, N. Antipa, W. Liberti, S. Dehaeck, K. Monakhova, F. L. Liu, K. Shen, R. Ng, and L. Waller, "Miniscope3d: optimized single-shot miniature 3d fluorescence microscopy," *Light: Sci. Appl.* **9**(1), 171 (2020).
22. F. L. Liu, G. Kuo, N. Antipa, K. Yanny, and L. Waller, "Fourier diffuserscope: Single-shot 3d fourier light field microscopy with a diffuser," arXiv preprint arXiv:2006.16343 (2020).
23. K. Monakhova, K. Yanny, N. Aggarwal, and L. Waller, "Spectral DiffuserCam: Lensless snapshot hyperspectral imaging with a spectral filter array," *Optica* **7**(10), 1298–1307 (2020).
24. N. Antipa, P. Oare, E. Bostan, R. Ng, and L. Waller, "Video from stills: Lensless imaging with rolling shutter," in *2019 IEEE International Conference on Computational Photography (ICCP)*, (IEEE, 2019), pp. 1–8.
25. C. Li, "An efficient algorithm for total variation regularization with applications to the single pixel camera and compressive sensing," Ph.D. thesis, Rice University (2010).
26. K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Trans. on Image Process.* **26**(9), 4509–4522 (2017).
27. E. Ryu, J. Liu, S. Wang, X. Chen, Z. Wang, and W. Yin, "Plug-and-play methods provably converge with properly trained denoisers," in *International Conference on Machine Learning*, (PMLR, 2019), pp. 5546–5557.
28. S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *2013 IEEE Global Conference on Signal and Information Processing*, (IEEE, 2013), pp. 945–948.
29. K. Zhang, Y. Li, W. Zuo, L. Zhang, L. Van Gool, and R. Timofte, "Plug-and-play image restoration with deep denoiser prior," arXiv preprint arXiv:2008.13751 (2020).
30. Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (red)," *SIAM J. Imaging Sci.* **10**(4), 1804–1844 (2017).
31. K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising with block-matching and 3d filtering," in *Image Processing: Algorithms and Systems, Neural Networks, and Machine Learning*, vol. 6064 (International Society for Optics and Photonics, 2006), p. 606414.
32. A. Sinha, J. Lee, S. Li, and G. Barbastathis, "Lensless computational imaging through deep learning," *Optica* **4**(9), 1117–1125 (2017).
33. M. T. McCann, K. H. Jin, and M. Unser, "Convolutional neural networks for inverse problems in imaging: A review," *IEEE Signal Process. Mag.* **34**(6), 85–95 (2017).

34. K. Monakhova, J. Yurtsever, G. Kuo, N. Antipa, K. Yanny, and L. Waller, "Learned reconstructions for practical mask-based lensless imaging," *Opt. Express* **27**(20), 28075–28090 (2019).
35. S. S. Khan, V. Sundar, V. Boominathan, A. N. Veeraraghavan, and K. Mitra, "Flatnet: Towards photorealistic scene reconstruction from lensless measurements," *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020).
36. X. Zhang, Q. Chen, R. Ng, and V. Koltun, "Zoom to learn, learn to zoom," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (2019), pp. 3762–3770.
37. D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (2018), pp. 9446–9454.
38. R. Heckel and P. Hand, "Deep decoder: Concise image representations from untrained non-convolutional networks," arXiv preprint arXiv:1810.03982 (2018).
39. G. Mataev, P. Milanfar, and M. Elad, "DeepRED: Deep image prior powered by RED," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, (2019), pp. 0–0.
40. J. Liu, Y. Sun, X. Xu, and U. S. Kamilov, "Image restoration using total variation regularized deep image prior," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, (IEEE, 2019), pp. 7715–7719.
41. E. Bostan, R. Heckel, M. Chen, M. Kellman, and L. Waller, "Deep phase decoder: self-calibrating phase microscopy with an untrained deep neural network," *Optica* **7**(6), 559–562 (2020).
42. K. Gong, C. Catana, J. Qi, and Q. Li, "Pet image reconstruction using deep image prior," *IEEE Trans. Med. Imaging* **38**(7), 1655–1665 (2019).
43. J. Cui, K. Gong, N. Guo, C. Wu, X. Meng, K. Kim, K. Zheng, Z. Wu, L. Fu, B. Xu, Z. Zhu, J. Tian, H. Liu, and Q. Li, "Pet image denoising using unsupervised deep learning," *Eur. J. Nucl. Med. Mol. Imaging* **46**(13), 2780–2789 (2019).
44. K. H. Jin, H. Gupta, J. Yerly, M. Stuber, and M. Unser, "Time-dependent deep image prior for dynamic MRI," arXiv preprint arXiv:1910.01684 (2019).
45. K. C. Zhou and R. Horstmeyer, "Diffraction tomography with a deep image prior," *Opt. Express* **28**(9), 12872–12896 (2020).
46. D. Van Veen, A. Jalal, M. Soltanolkotabi, E. Price, S. Vishwanath, and A. G. Dimakis, "Compressed sensing with deep image prior and learned regularization," arXiv preprint arXiv:1806.06438 (2018).
47. G. Kuo, N. Antipa, R. Ng, and L. Waller, "DiffuserCam: diffuser-based lensless cameras," in *Computational Optical Sensing and Imaging*, (Optical Society of America, 2017), pp. CTu3B–2.
48. A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imaging Sci.* **2**(1), 183–202 (2009).
49. S. Boyd, N. Parikh, and E. Chu, *Distributed optimization and statistical learning via the alternating direction method of multipliers* (Now Publishers Inc, 2011).
50. U. S. Kamilov, "A parallel proximal algorithm for anisotropic total variation minimization," *IEEE Trans. on Image Process.* **26**(2), 539–548 (2017).
51. M. Maggioni, V. Katkovnik, K. Egiazarian, and A. Foi, "Nonlocal transform-domain filter for volumetric data denoising and reconstruction," *IEEE Trans. on Image Process.* **22**(1), 119–133 (2013).
52. Luminet, "Technical Data and Downloads," <http://www.luminetco.com/downloads/data-sheets> (2017). [Online; accessed 19-July-2008].
53. R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, (2018), pp. 586–595.
54. Z. Wang, E. P. Simoncelli, and A. C. Bovik, "Multiscale structural similarity for image quality assessment," in *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, vol. 2 (Ieee, 2003), pp. 1398–1402.
55. R. Ennis, F. Schiller, M. Toscani, and K. R. Gegenfurtner, "Hyperspectral database of fruits and vegetables," *J. Opt. Soc. Am. A* **35**(4), B256–B266 (2018).
56. K. Monakhova, V. Tran, G. Kuo, and L. Waller, "Untrained deep network (UDN) reconstructions for compressive lensless photography," GitHub (2021) <https://github.com/Waller-Lab/UDN/>