

UC Riverside

UC Riverside Previously Published Works

Title

Joint Image and Depth Estimation With Mask-Based Lensless Cameras.

Permalink

<https://escholarship.org/uc/item/49r3b970>

Authors

Zheng, Yucheng
Asif, M Salman

Publication Date

2020

DOI

10.1109/TCI.2020.3010360

Peer reviewed

Joint Image and Depth Estimation with Mask-Based Lensless Cameras

Yucheng Zheng and M. Salman Asif

Abstract—Mask-based lensless cameras replace the lens of a conventional camera with a custom mask. These cameras can potentially be very thin and even flexible. Recently, it has been demonstrated that such mask-based cameras can recover light intensity and depth information of a scene. Existing depth recovery algorithms either assume that the scene consists of a small number of depth planes or solve a sparse recovery problem over a large 3D volume. Both these approaches fail to recover the scenes with large depth variations. In this paper, we propose a new approach for depth estimation based on an alternating gradient descent algorithm that jointly estimates a continuous depth map and light distribution of the unknown scene from its lensless measurements. We present simulation results on image and depth reconstruction for a variety of 3D test scenes. A comparison between the proposed algorithm and other method shows that our algorithm is more robust for natural scenes with a large range of depths. We built a prototype lensless camera and present experimental results for reconstruction of intensity and depth maps of different real objects.

Index Terms—Lensless imaging, flatcam, depth estimation, non-convex optimization, alternating minimization.

I. INTRODUCTION

Depth estimation is an important and challenging problem that arises in a variety of applications including computer vision, robotics, and autonomous systems. Existing depth estimation systems use stereo pairs of conventional (lens-based) cameras or time-of-flight sensors [2]–[4]. These cameras can be heavy, bulky, and require large space for their installation. Therefore, their adoption in portable and lightweight devices with strict physical constraints is still limited.

In this paper, we propose a joint image and depth estimation framework for a computational lensless camera that consists of a fixed, binary mask placed on top of a bare sensor. Such mask-based cameras offer an alternative design for building cameras without lenses. A recent example of a mask-based lensless camera is known as FlatCam [5]. In contrast with a lens-based camera that is designed to map every point in the scene to a single pixel on the sensor, every sensor pixel in a FlatCam records light from many points in the scene. A single point source in the scene casts a shadow of the mask on the sensor, which shifts if the point moves parallel to the sensor plane and expands/shrinks if the point source moves toward/away from the sensor plane. The measurements recorded on the sensor thus represent a superposition of shifted and scaled versions

of the mask shadows corresponding to light sources in different directions and depths. Image and depth information about the scene is thus encoded in the measurements, and we can solve an inverse problem to estimate both of them.

To jointly estimate the depth and light distribution, we propose a two-step approach that consists of an initialization step and an alternating gradient descent step to minimize our objective. To preserve sharp edges in the image intensity and depth map, we include an adaptive regularization penalty in our objective function.

An overview of the reconstruction framework is illustrated in Figure 1. In this paper, we use the same sampling framework proposed in [6]. We initialize the estimate of the depth map by selecting a single plane or solving the greedy algorithm proposed in [6]. The greedy method assumes that the scene consists of a small number of depth planes and fails to recover scene with continuous depth variations. The method proposed in this paper can estimate continuous depth by minimizing an objective function with respect to image intensity and depth via alternating gradient descent. We present extensive simulation and real experimental results with different objects.

The main contributions of this paper are as follows.

- We propose a new computational framework for joint estimation of light intensity and depth maps from a single image of a mask-based lensless camera. In contrast to other methods, our method estimates the depth map on a continuous domain. Our algorithm consists of a careful initialization step based on greedy pursuit and an alternating minimization step based on gradient descent.
- The problem of joint image and depth recovery is highly nonconvex. To tackle this issue, we present different regularization schemes that offer robust recovery on a diverse dataset.
- We present simulation results on standard 3D datasets and demonstrated a significant improvement over existing methods for 3D imaging using coded mask-based lensless cameras.
- We built a hardware prototype to capture measurements of real objects. We present image and depth reconstruction results of these real objects using our proposed algorithm and a comparison with existing methods.

II. RELATED WORK

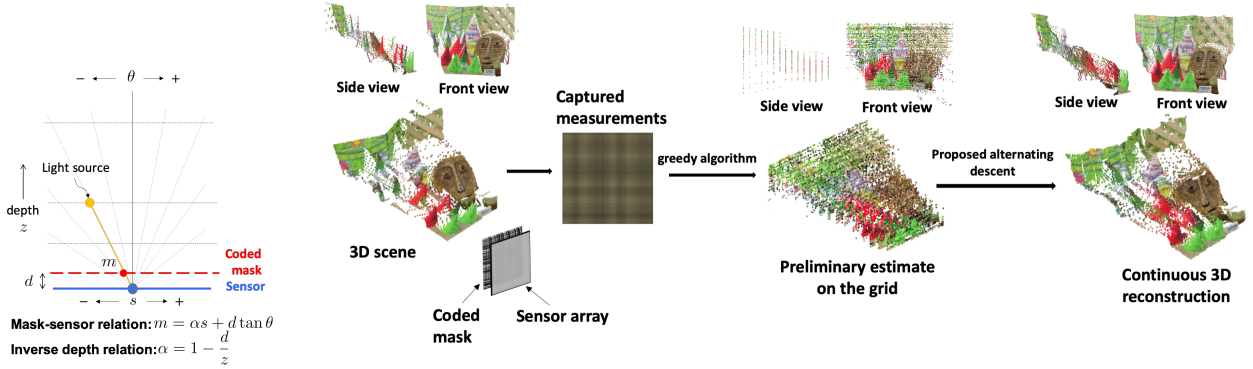
A pinhole camera, also known as *camera obscura*, is the simplest example of a mask-based lensless camera. Even though a pinhole can easily provide an image of the scene onto a sensor plane, the image quality is often severely affected

Y. Zheng and M. Asif are with the Department of Electrical and Computer Engineering, University of California, Riverside, CA, 92521 USA (e-mail: yzhen069@ucr.edu; sasif@ece.ucr.edu).

This work was supported in part by Google Faculty Award.

A shorter version with preliminary results was presented in [1]

Additional results are available in the supplementary material.



(a) 1D imaging model for a planar sensor with a coded mask placed at distance d . Light rays from a light source at location (θ, z) are received by all the sensor pixels. A light ray that hits sensor pixel s passes through mask at location m .

(b) An overview of the proposed intensity and depth estimation framework. Consider a natural scene as a 3D point cloud, where each point represents a light source located at a different depth. The camera consists of a fixed, coded mask placed on top of an image sensor. Every point in the scene casts a shadow of the mask on the sensor plane. Each sensor pixels records a linear combination of the scene modulated by the mask pattern. The recovery algorithm consists of two steps. (1) Initialization using a greedy depth selection method. (2) An alternating gradient descent-based refinement algorithm that jointly estimates the light distribution and depth map on a continuous domain.

Fig. 1: A coded mask-based imaging model and an overview of the proposed continuous depth estimation framework.

by noise because the amount of light collected is limited by the pinhole aperture [7]. Coded aperture-based lensless cameras avoid this problem by increasing the number of pinholes and allowing more light to reach the sensor [5], [8]–[12]. In contrast to a pinhole camera where only one inverted image of the scene is obtained through a single pinhole, the measurements captured through a coded-mask are a linear combination of all the pinhole images under every mask element. To recover an image of the scene, we need to solve a computational image recovery problem [5], [8], [12].

Recent work on mask-based lensless imaging broadly falls into two categories. FlatCam [6] uses a separable mask aligned with the sensor such that the sensor measurements corresponding to a plane at a fixed depth from the sensor can be written as a separable system. DiffuserCam [12] assumes that the mask size and angular span of the object are small enough so that the sensor measurements of a plane can be modeled as a convolution of the mask pattern with image intensity at that plane. The convolutional model can be computationally efficient if the object falls within a small angular range because we can use fast Fourier transform to compute convolutions. The separable model does not require a small angular range assumption. A number of methods based on deep learning have also been developed recently for both separable and convolutional imaging models to recover images at a fixed depth plane [13]–[15].

A coded aperture system offers another advantage by encoding light from different directions and depths differently. The depth-dependent imaging capability in coded aperture systems is known since the pioneering work in this domain [8], [16]. However, the classical methods usually assume that the scene consists of a single plane at known depth. In this paper, we assume that the depth map is arbitrarily distributed on a continuous domain and the true depth map is unknown at the time of reconstruction.

The 3D lensless imaging problem has also recently been

studied in [6], [11], [12], [17], [18]. These methods can broadly be divided into two categories. In the first category, the 3D scene is divided into a finite number of voxels. To recover the 3D light distribution, these methods solve an ℓ_1 norm-based recovery problem under the assumption that the scene is very sparse [12], [17]. In the second category, the 3D scene is divided into an intensity map and multiple depth planes such that each pixel is assigned one intensity and depth. To solve the intensity and depth recovery problem, these methods either sweep through the depth planes [11], [18] or assign depth to each pixel using a greedy method [6]. Our proposed method belongs to the second category in which we model the image intensity and depth separately and assume that the depth values of the scene are distributed on a continuous domain. To recover the 3D scene, we jointly estimate the image intensity and depth map from the available sensor measurements.

Joint estimation of image intensity and depth map can be viewed as a nonlinear inverse problem in which the sampling function is dependent on scene depth. Similar inverse problem also arises in many other fields such as direction-of-arrival estimation in radar [19], super-resolution [20] and compressed sensing [21]–[23]. Similar to the joint estimation of image intensity and depth, the solution approaches to these problems consists of two main steps: identification of signal bases and the estimation of signal intensities based on the identified bases. The problem of identifying the signal bases from continuously varying candidates is often called off-the-grid signal recovery. The methods for solving the off-the-grid signal recovery problems can be divided into two main types. The first approach formulates the problem as a convex program on a continuous domain and solves it using an atomic norm minimization approach [24], [25]. The second approach linearizes the problem for the optimization parameter using a first-order approximation at every iteration [20], [26]. Our proposed algorithm is inspired by the second approach.

Mask-based lensless cameras have traditionally been used

for imaging light at wavelengths beyond the visible spectrum [9], [10]. Other examples related to mask-based cameras include controllable aperture and employing coded-mask for compressed sensing and computational imaging [27], [28], distributed lensless camera [29], single pixel camera [30] and external mask setting [31].

Coded masks have also recently been used with conventional lens-based cameras to estimate depth and lightfield [32]–[35]. Recently, a number of data-driven methods have been proposed to design custom phase masks and optical elements to estimate depth from a single image [36], [37]. An all-optical diffractive deep neural network is proposed in [38], [39], which can perform pattern recognition tasks such as handwritten digits classification using optical mask layers. Such networks can literally process images at a lightning-fast pace with near-zero energy cost.

III. METHODS

A. Imaging Model

We divide the 3D scene under observation into $N \times N$ uniformly spaced directions. We use θ_i and θ_j to denote the angular directions of a light source with respect to the center of the sensor. The intensity and depth of the light source are denoted using $l_{i,j}$ and $z_{i,j}$ respectively. Figure 1(a) depicts the geometry of such an imaging model. A planar coded-mask is placed on top of a planar sensor array at distance d . The $M \times M$ sensor array captures lights coming from the scene modulated by the coded-mask.

Every light source in the scene casts a shadow of the mask on the sensor array, which we denote using basis functions ψ . We use s_u and s_v to index a pixel on the rectangular sensor array. The shadow cast by a light source with unit intensity at $(\theta_i, \theta_j, z_{i,j})$ can be represented as the following basis or point spread function:

$$\psi_{i,j}(s_u, s_v) = \text{mask}[\alpha_{i,j}s_u + d \tan(\theta_i), \alpha_{i,j}s_v + d \tan(\theta_j)], \quad (1)$$

where $\text{mask}[u, v]$ denotes the transmittance of the mask pattern at location (u, v) on the mask plane and $\alpha_{i,j}$ is a variable that is related to the physical depth $z_{i,j}$ with the following inverse relation:

$$\alpha_{i,j} = 1 - \frac{d}{z_{i,j}}, \quad (2)$$

If the 3D scene consists of only a single point source at (θ_i, θ_j) with light intensity $l_{i,j}$, the measurement captured at sensor pixel (s_u, s_v) would be

$$y(s_u, s_v) = \psi_{i,j}(s_u, s_v)l_{i,j}. \quad (3)$$

The measurement recorded on any sensor pixel is the summation of contributions from each of the point sources in the 3D scene. The imaging model for a single sensor pixel can be represented by

$$y(s_u, s_v) = \sum_{i=1}^N \sum_{j=1}^N \psi_{i,j}(s_u, s_v)l_{i,j}. \quad (4)$$

We can write the imaging model for the entire sensor in a compact form as

$$\mathbf{y} = \Psi(\alpha)\mathbf{l} + e, \quad (5)$$

where $\mathbf{y} \in \mathbb{R}^{M^2}$ is a vectorized form of an $M \times M$ matrix that denotes sensor measurements, $\mathbf{l} \in \mathbb{R}^{N^2}$ is a vectorized form of an $N \times N$ matrix that denotes light intensity from all the locations $(\theta_i, \theta_j, \alpha_{i,j})$, and Ψ is a matrix with all the basis functions corresponding to $\theta_i, \theta_j, \alpha_{i,j}$. The basis functions in (5) are parameterized by the unknown $\alpha \in \mathbb{R}^{N^2}$ and e denotes noise and other nonidealities in the system.

We can jointly estimate light distribution (1) and inverse depth map $(\alpha)^1$ using the following optimization problem:

$$\underset{\alpha, \mathbf{l}}{\text{minimize}} \quad \frac{1}{2} \|\mathbf{y} - \Psi(\alpha)\mathbf{l}\|_2^2. \quad (6)$$

Note that if we know the true values of α (or we fix it to something), then the problem in (6) reduces to a linear least-squares problem that can be efficiently solved via standard solvers. On the other hand, if we fix the value of \mathbf{l} , the problem remains nonlinear with respect to α . In the next few sections we discuss our approach for solving the problem in (6) via alternating minimization.

B. Initialization

Since the minimization problem in (6) is not convex, a proper initialization is often needed to ensure convergence to a local minimum close to the optimal point. A naïve approach is to initialize all the point sources in the scene at the same depth plane. To select an initial depth plane, we sweep through a set of candidate depth planes and perform image reconstruction on one depth plane at a time by solving the following linear least squares problem:

$$\underset{\mathbf{l}}{\text{minimize}} \quad \frac{1}{2} \|\mathbf{y} - \Psi(\alpha)\mathbf{l}\|_2^2. \quad (7)$$

We evaluate the loss value for all the candidate depth planes and picked the one with the smallest loss as our initialized depth. The mask basis function in (1) changes as we change α , which has an inverse relation with the scene depth. We select candidate depth corresponding to uniformly sampled values of α , which yields non-uniform sampling of the physical scene depth. The single-depth initialization approach is computationally simple and provides a reasonable initialization of light distribution to start with, especially when the scene is far from the sensor.

Our second approach for initialization is the greedy method proposed in [6]. Greedy algorithms are widely used for sparse signal recovery [21]–[23]. Based on these algorithms, [6] proposed a greedy depth pursuit algorithm for depth estimation from FlatCam [5]. The algorithm works by iteratively updating the depth surface that matches the observed measurements the best.

The depth pursuit method assumes that the scene consists of a small number of predefined depth planes. We start the program by initializing all the pixels at a single depth plane

¹ α has an inverse relation with the depth map (2); therefore we refer to it as inverse depth map throughout the paper.

and the estimation of light intensities \mathbf{I} based on the initialized depth map. The first step is to select new candidate values for α . The new candidates are selected using the basis vectors that are mostly correlated with the current residual of the estimate. In the second step, new candidates for α are appended to the current estimate. We solve a least squares problem using the appended α . In the third step, we prune the α by selecting $\alpha_{i,j}$ as the value corresponding to the largest magnitude of $l_{i,j}$. Although this method may not estimate the off-grid point sources well, it produces a good preliminary estimate of the scene.

C. Refinement via Alternating Gradient Descent

To solve the minimization problem in (6), we start with the preliminary image and depth estimates from the initialization step and alternately update depth and light distribution via gradient descent. The main computational task in gradient descent method is computing the gradient of the loss function w.r.t. α . To compute that gradient, we expand the loss function in (6) as

$$L = \frac{1}{2} \sum_{u,v=1}^M (y(s_u, s_v) - \sum_{i,j=1}^N \psi_{i,j}(s_u, s_v) l_{i,j})^2 \quad (8)$$

We define $R_{u,v} = y(s_u, s_v) - \sum_{i,j=1}^N \psi_{i,j}(s_u, s_v) l_{i,j}$ as the residual approximation error at location (s_u, s_v) . The derivatives of the loss function with respect to the $\alpha_{i,j}$ is given as

$$\frac{\partial L}{\partial \alpha_{i,j}} = \sum_{u,v=1}^M R_{u,v} \frac{\partial R_{u,v}}{\partial \alpha_{i,j}} = -l_{i,j} \sum_{u,v=1}^M R_{u,v} \frac{\partial \psi_{i,j}(s_u, s_v)}{\partial \alpha_{i,j}}. \quad (9)$$

We compute the derivatives of sensor value with respect to the $\alpha_{i,j}$ using the total derivative² as follows.

$$\begin{aligned} \frac{\partial \psi_{i,j}(s_u, s_v)}{\partial \alpha_{i,j}} &= \frac{\partial \psi_{i,j}(s_u, s_v)}{\partial u_{i,j}} \frac{\partial u_{i,j}}{\partial \alpha_{i,j}} + \frac{\partial \psi_{i,j}(s_u, s_v)}{\partial v_{i,j}} \frac{\partial v_{i,j}}{\partial \alpha_{i,j}} \\ &= \frac{\partial \psi_{i,j}(s_u, s_v)}{\partial u_{i,j}} s_u + \frac{\partial \psi_{i,j}(s_u, s_v)}{\partial v_{i,j}} s_v. \end{aligned} \quad (10)$$

$u_{i,j} = \alpha_{i,j} s_u + d \tan(\theta_i)$ and $v_{i,j} = \alpha_{i,j} s_v + d \tan(\theta_j)$ denote two dummy variables that also correspond to the specific location on the mask where a light ray from a point source at angle (θ_i, θ_j) and depth $\alpha_{i,j}$ and sensor pixel at (s_u, s_v) intersects with the mask plane. The terms in $\frac{\partial \psi_{i,j}(s_u, s_v)}{\partial u_{i,j}}$, $\frac{\partial \psi_{i,j}(s_u, s_v)}{\partial v_{i,j}}$ can be viewed as the derivatives of mask pattern along the respective spatial coordinates and evaluated at $u_{i,j}, v_{i,j}$. We compute these derivatives using finite-difference of $\psi_{i,j}(s_u, s_v)$ over a fine grid and linear interpolation.

D. Algorithm Analysis

To solve the non-linear least squares problem in (6) in our algorithms, we compute the gradient derived in (10) and use

²Recall that the total derivative of a multivariate function $f(x, y)$ is $\frac{\partial f(x, y)}{\partial x} dx + \frac{\partial f(x, y)}{\partial y} dy$.

it as input of a optimization solver. Suppose ψ_i and ψ_j denote the basis function vectors evaluated on a 1D mask as

$$\begin{aligned} \psi_i(s_u) &= \text{mask} [\alpha_{i,j} s_u + d \tan(\theta_i)] \\ \psi_j(s_v) &= \text{mask} [\alpha_{i,j} s_v + d \tan(\theta_j)]. \end{aligned} \quad (11)$$

If we use a separable mask pattern, then the 2D mask function $\psi_{i,j}$ in (1) can be computed as the outer product of two vectors given as $\psi_{i,j} = \psi_i \psi_j^T$. Similarly, we define 1D sub-gradient function g as

$$\begin{aligned} g_i(s_u) &= \frac{\partial \psi_{i,j}(s_u, s_v)}{\partial u_{i,j}} \\ g_j(s_v) &= \frac{\partial \psi_{i,j}(s_u, s_v)}{\partial v_{i,j}}, \end{aligned} \quad (12)$$

Similar to (10), the functions $\frac{\partial \psi_{i,j}(s_u, s_v)}{\partial u_{i,j}}$ and $\frac{\partial \psi_{i,j}(s_u, s_v)}{\partial v_{i,j}}$ are the sub-gradient functions along the 1D mask. It takes non-negative values at locations where mask pattern value changes and takes zero value at the other places. Using the derivation in (10), the matrix contains $\frac{\partial \psi_{i,j}(s_u, s_v)}{\partial \alpha_{i,j}}$ at all (s_u, s_v) can be computed using the following sum of two vector outer products.

$$\frac{\partial \psi_{i,j}}{\partial \alpha_{i,j}} = g_i \psi_j^T + \psi_i g_j^T \quad (13)$$

Using the derivations in (9), the derivative of loss function with respect to depth value can be computed using the following matrix multiplications, where R refers to the matrix of residual $R_{u,v}$ at all (s_u, s_v)

$$\frac{\partial L}{\partial \alpha_{i,j}} = g_i^T R \psi_j + \psi_i^T R g_j \quad (14)$$

Suppose we have $M \times M$ pixels on sensor array. The computation in (14) takes $2M^2 + 2M$ multiplications. We then feed our gradients to `minfunc` solver [40] with L-BFGS algorithm [41] to solve the non-linear optimization problem in (6).

E. Regularization Approaches

ℓ_2 regularization on spatial gradients. The optimization problem in (6) is highly non-convex and contains several local minima; therefore, the estimate often gets stuck in some local minima and the estimated intensity and depth maps are coarse. To improve the performance of our algorithm for solving the non-convex problem in (6), we seek to exploit additional structures in the scene. A standard assumption is that the depth of neighboring pixels is usually close, which implies that the spatial differences of (inverse) depth map are small. To incorporate this assumption in our model, we add a quadratic regularization term on the spatial gradients of the inverse depth map to our loss function. The quadratic regularization term is defined on an $N \times N$ inverse depth map matrix α and can be written as

$$\begin{aligned} R(\alpha) &= \sum_{i,j=1}^N (\alpha_{i,j} - \alpha_{i+1,j})^2 + (\alpha_{i,j} - \alpha_{i,j+1})^2 \\ &= \|\nabla_r \alpha\|_F^2 + \|\nabla_c \alpha\|_F^2, \end{aligned} \quad (15)$$

where the operators ∇_r, ∇_c compute spatial differences along rows and columns, respectively. We call this regularization an

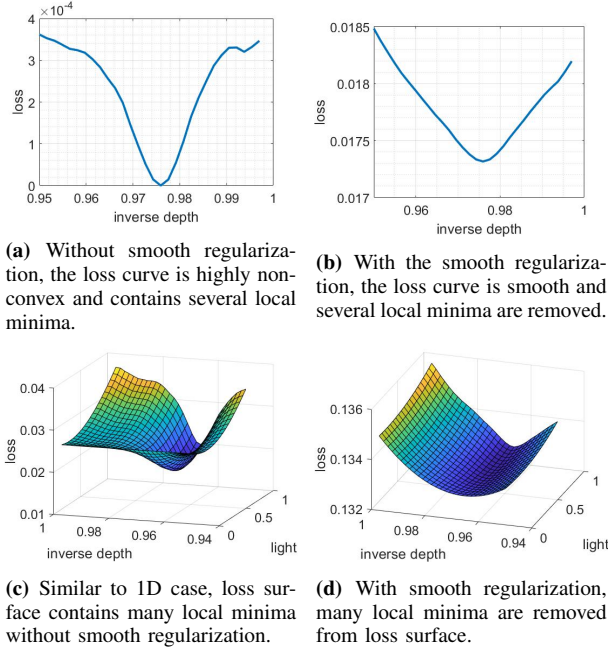


Fig. 2: A comparison between objective loss functions without and with smooth regularization. The inverse depth axis refers to the value of α .

ℓ_2 norm-based total variation (TV- ℓ_2) in this paper. Figure 2 illustrates the effect of the depth regularization. From Figure 2, we observe that smoothness regularization improves the loss function by removing several local minima. We also observed this effect in our simulations for a high-dimensional depth recovery problem, which is not very sensitive to initialization with depth regularization.

Weighted ℓ_2 regularization on spatial gradients. Even though smoothness regularization on the inverse depth map removes some local minima and helps with converge, it does not respect the sharp edges in the depth map. To preserve sharp discontinuities in the (inverse) depth map, we used the following adaptive weighted regularization inspired from [42]:

$$R_W(\alpha) = \sum_{i,j=1}^N W_{i,j}^c (\alpha_{i,j} - \alpha_{i+1,j})^2 + W_{i,j}^r (\alpha_{i,j} - \alpha_{i,j+1})^2, \quad (16)$$

where $W_{i,j}^{r,\alpha}$ and $W_{i,j}^{c,\alpha}$ denote weights for row and column differences, respectively. We aim to select these weights to promote depth similarity for neighboring pixels, but avoid smoothing the sharp edges. To promote this, we selected weights with exponential decay in our experiments that we compute as

$$W_{i,j}^r = \exp\left(-\frac{(\alpha_{i,j} - \alpha_{i+1,j})^2}{\sigma}\right)$$

$$W_{i,j}^c = \exp\left(-\frac{(\alpha_{i,j} - \alpha_{i,j+1})^2}{\sigma}\right). \quad (17)$$

Such a weighted regularization forces pixels that have depth within a small range of one another to be smooth and does not penalize the points that have larger gap in depth (which indicates the presence of an edge). This helps preserve sharp

edges in the reconstructed depth estimates. This weighting approach is analogous to bilateral filtering approach for image denoising [43], [44].

The regularized estimation problem for image and depth can be written in the following form:

$$\underset{\alpha, \mathbf{l}}{\text{minimize}} \quad \frac{1}{2} \|\mathbf{y} - \Psi(\alpha)\mathbf{l}\|_2^2 + \lambda R_W(\alpha). \quad (18)$$

We call this regularization approach weighted TV- ℓ_2 and solve it by alternately updating the inverse depth map α and light intensity \mathbf{l} . A pseudocode of the algorithm is presented at Algorithm 1.

Algorithm 1 Weighted TV- ℓ_2 regularized optimization

Input: Sensor measurements: \mathbf{y}

Output: Light distribution and inverse depth map: \mathbf{l}, α

Initialization via greedy algorithm:

Compute α and \mathbf{l} with depth pursuit algorithm in [6].

Refinement via alternating gradient descent:

for $k = 1 : k_{\max}$ **do**

$$\hat{\alpha}^k = \underset{\alpha}{\text{argmin}} \quad \frac{1}{2} \|\mathbf{y} - \Psi(\alpha)\mathbf{l}^{k-1}\|_2^2 + \lambda R_W(\alpha)$$

$$\hat{\mathbf{l}}^k = \underset{\mathbf{l}}{\text{argmin}} \quad \frac{1}{2} \|\mathbf{y} - \Psi(\hat{\alpha}^k)\mathbf{l}\|_2^2$$

end for

return $\hat{\mathbf{l}}$ and $\hat{\alpha}$

ℓ_1 regularization on spatial gradients. It is well-known that the ℓ_1 norm regularization enforces the solution to be sparse. We add an ℓ_1 -based total variation norm [45] of the depth to our optimization problem. By enforcing the sparsity of spatial gradients, the edges of (inverse) depth map can be preserved. The ℓ_1 norm-based TV regularization term is given as

$$R_{TV}(\alpha) = \sum_{i,j=1}^N |\alpha_{i,j} - \alpha_{i+1,j}| + |\alpha_{i,j} - \alpha_{i,j+1}|$$

$$= \|\nabla_r \alpha\|_1 + \|\nabla_c \alpha\|_1. \quad (19)$$

To solve the nonlinear optimization problem with ℓ_1 norm regularization, we write the optimization problem as

$$\underset{\alpha, \mathbf{l}}{\text{minimize}} \quad \frac{1}{2} \|\mathbf{y} - \Psi(\alpha)\mathbf{l}\|_2^2 + \lambda (\|\mathbf{d}_r\|_1 + \|\mathbf{d}_c\|_1)$$

$$\text{s.t. } \mathbf{d}_r = \nabla_r \alpha, \quad \mathbf{d}_c = \nabla_c \alpha. \quad (20)$$

We solve this problem (20) using a split-Bregman method [46].

F. Computational complexity

The main computational and storage cost of the proposed method arises from the forward and adjoint of the imaging operators. Since the depth of any pixel can have an arbitrary value in a continuous domain, we cannot precompute the imaging operators. At every iteration, we first interpolate the separable mask patterns according to the current estimate of α . The time complexity of forward imaging operator as given in (5) will be $O(M^2 N^2)$ because we have to add up depth-dependent response corresponding to every angle.

IV. SIMULATION RESULTS

In this section, we present simulation results to evaluate the performance of our methods under different noise levels and sensor sizes. We also present a comparison of our proposed method with two existing methods for 3D imaging with lensless cameras. Additional experiments on the reconstruction of a single depth plane and the effect of numbers of sensor pixels on the reconstruction are included in the Supplementary material.

A. Simulation Setup

To validate the performance of the proposed algorithm, we simulate a lensless imaging system using a binary planar mask with a separable maximum length sequence (MLS) pattern [47] that is placed 4mm away from a planar sensor array. We used an MLS sequence of length 1024 and converted all the -1 s to 0s to create a separable binary pattern. We used square mask features, each of which is $30\mu\text{m}$ wide. Since we optimize the objective function in (6) with respect to α and need to compute the gradient in (9), we require the mask function to be smooth and differentiable with respect to α . Therefore, we convolved the binary pattern with a Gaussian blur kernel of length $15\mu\text{m}$ and standard deviation 5. In our simulations, we do not explicitly model the diffraction blur. However, the Gaussian blur kernel that we apply to the mask function can be viewed as an approximation of the diffraction blur. The sensor contains 512×512 square pixels, each of which is $50\mu\text{m}$ wide. The chief ray angle of each sensor pixel is $\pm 18^\circ$. We assume that there is no noise added to the sensor measurements. In our experiments for continuous depth estimation, we fixed all the parameters to these default values and analyze the performance with respect to a single parameter.

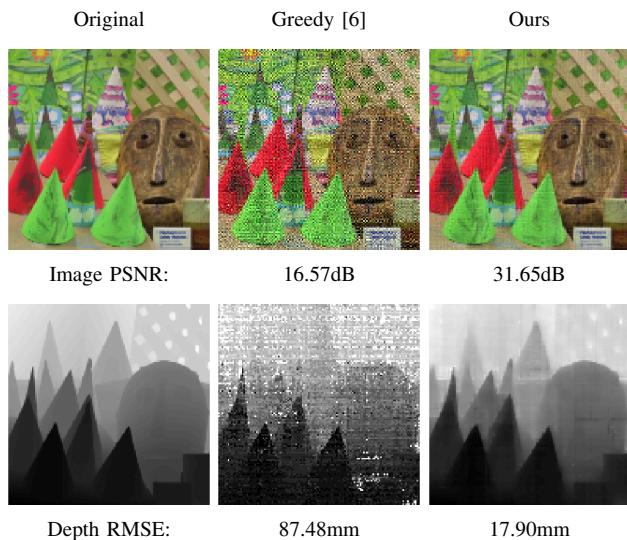


Fig. 3: Left to right: original image and depth of the Cones scene; image and depth initialized via greedy algorithm [6]; depth estimation using weighted ℓ_2 -based regularization. The depth in this scene varies from around 0.99m to 1.7m.

B. Reconstruction of Scenes with Continuous Depth

Depth datasets: We performed all our experiments on 3D images created using light intensities and depth information from Middlebury [48], Make3D [49], [50] and NYU Depth [51], the test scenes and their depth ranges are listed in Table I.

Test datasets	Min depth (m)	Max depth (m)
Sword	0.65	0.95
Playable	1.47	3.75
Cones	0.99	1.70
Corner	3.93	10.60
Whiteboard	1.08	2.90
Playroom	1.62	2.93
Moebius	0.74	1.23
Books	0.73	1.27

TABLE I: Analysis experiments are performed on multiple scenes picked from Middlebury [48], Make3D [49], [50] and NYU Depth [51]. Results of the two scenes above line are presented within the main text, while the rest of them are reported in the supplementary material.

Initialization via greedy method: Let us further discuss our simulation setup using the *Cones* scene, for which the results are presented in Figure 3. We simulated the 3D scene using depth data from Middlebury dataset [48]. We sample the scene at uniform angles to create a 128×128 image and its (inverse) depth map with the same size. We can compute the physical depth from α using (2). In our simulation, the depth of this scene ranges from around 0.99m to 1.7m. We used depth pursuit greedy algorithm in [6] as our initialization method. We selected 15 candidate depths by uniformly sampling the inverse depth values α from 0.996 to 0.9976, which gives an effective depth in the same range as the original depth. Since we are trying to gauge the performance for off-the-grid estimate of depth, the candidate values of α are not exactly the same as the true values of α in our simulations. The output of the initialization algorithm is then fed into the alternating gradient descent method.

Performance metrics: We evaluate the performance of recovered image intensity and depth independent of each other. We report the peak signal to noise ratio (PSNR) of the estimated intensity distribution and root mean squared error (RMSE) of the estimated depth maps for all our experiments. The estimates for image intensity and depth maps for the initialization and our proposed weighted $\text{TV-}\ell_2$ method are shown in Figure 3, along with the PSNR and RMSE. We can observe that both image and depth estimation from greedy method [6] contain several spikes because of the model mismatch with the predefined depth grid. In contrast, many of these spikes are removed in the estimations from the proposed algorithm with weighted $\text{TV-}\ell_2$ while the edges are preserved.

Comparison of regularization methods: Here we present a comparison between three different regularization approaches. We reconstruct image intensity and (inverse) depth map using the same measurements with $\text{TV-}\ell_2$, weighted $\text{TV-}\ell_2$, and $\text{TV-}\ell_1$ regularization. The results are shown in Figure 4. Compared to the $\text{TV-}\ell_2$ method, we observe that both weighted

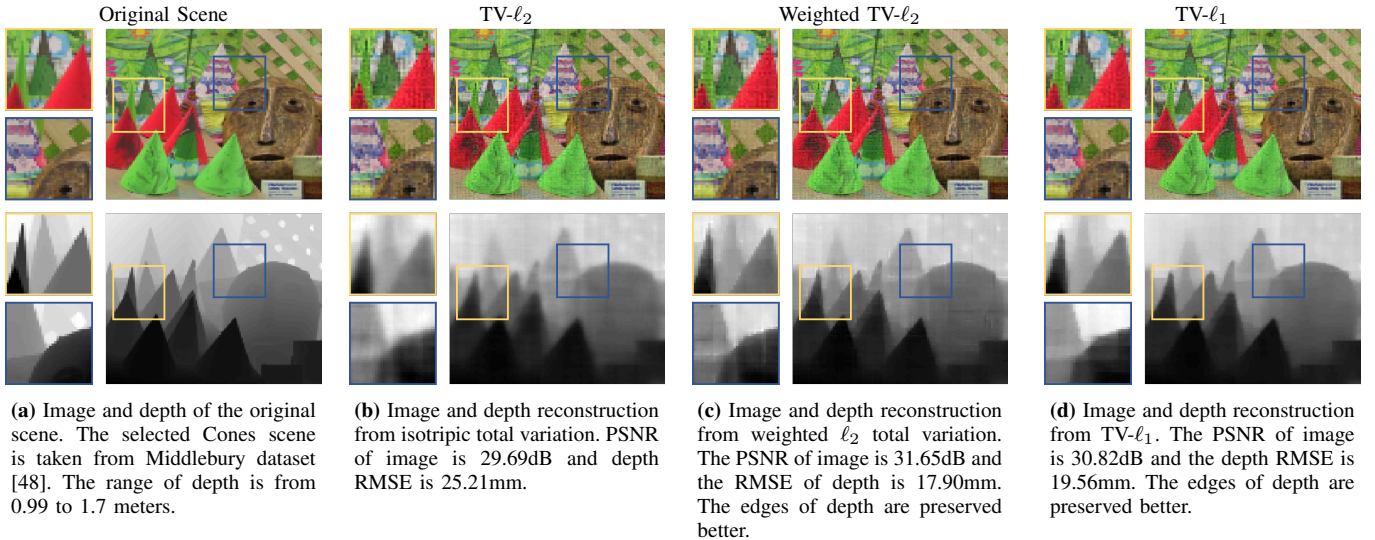


Fig. 4: Comparison between reconstructions using three different regularization approaches from the same measurements.

$TV-\ell_2$ and $TV-\ell_1$ preserve the sharp edges in image and depth estimates. Overall, in our experiments, weighted $TV-\ell_2$ provided the best results. Therefore, we used that as our default method for the rest of the paper.

C. Effects of Noise

Sensor noise exists widely in any observation process. The amplitude of noise depends on the intensities of sensor measurements and can adversely affect the reconstruction results. To investigate the effect of noise on our algorithm, we present simulation results for the reconstruction of scenes from the same sensor measurements under different levels of additive white Gaussian noise. The experiments are performed on multiple 3D scenes listed in Table I. Some examples of reconstruction with different levels of noise are shown in Figure 5.

The plots recording PSNR of image intensities and RMSE of depth maps over a range of measurement SNR values are presented in Figure 6. As we can observe from the curves that the quality of both estimated image and depth improve when the measurements have small noise (high SNR) and the quality degrades as we add more noise in the measurements (low SNR). Another observation we can make is that the scenes that are farther away have higher RMSE. This aspect is understandable because as the scenes move farther, α of the scene pixels all get very close to 1 and we cannot resolve fine depth variations in the scene.

D. Size of Sensor

In conventional disparity-based depth estimation method [2], the quality of reconstructed depth depends on the disparity between frames captured from multiple camera views. Larger distance between camera viewing positions results in better depth estimation accuracy. In a lensless imaging system, we can think of each pinhole on the mask and the sensor area behind the mask as a tiny pinhole camera. The analogy only

goes this far, because we do not record images from these tiny pinhole cameras separately; instead, we record a multiplexed version of all the views. The disparity between different points on the sensors, however, does affect our ability to resolve the depth of the scene, which is determined by the size of sensor.

To analyze the effect of disparity in our system, we performed experiments with three different sizes of sensor pixels from $25\mu\text{m}$, $50\mu\text{m}$, and $100\mu\text{m}$. For comparison, the number of sensor pixels and other parameters are set to the default settings as described earlier. No noise is included in this experiment. Results in terms of reconstructed image and depth maps are presented in Figure 7, where we observe that the quality of depth reconstruction improves as we increase the size of sensor pixels. The results in Figure 7 demonstrate that increasing the disparity of viewing points increases the depth reconstruction quality.

E. Comparison with Existing Methods

Finally, we present a comparison of our proposed algorithm and two other methods for 3D recovery with lensless cameras. In our method, we estimate light intensity and a depth map over continuous domain. The greedy method in [6] also estimates intensity and depth separately, but the depth map for any angle is restricted to one of the predetermined planes. Three-dimensional recovery using lensless cameras for 3D fluorescence microscopy was presented in [12] and [17], which estimate the entire 3D volume of the scene sampled over a predetermined 3D grid. Since the unknown volume scene in microscopy is often very sparse, the 3D scene recovery problem is solved as a sparse recovery problem for the light intensity over all the grid voxels. The result is a light distribution over the entire 3D space. We call this method 3D Grid and use the code provided in [12] to solve the 3D recovery problem using the forward model and measurements from our simulation setup.

The imaging experiments in [12] and [17] are aimed at fluorescence imaging in which objects are mostly transparent

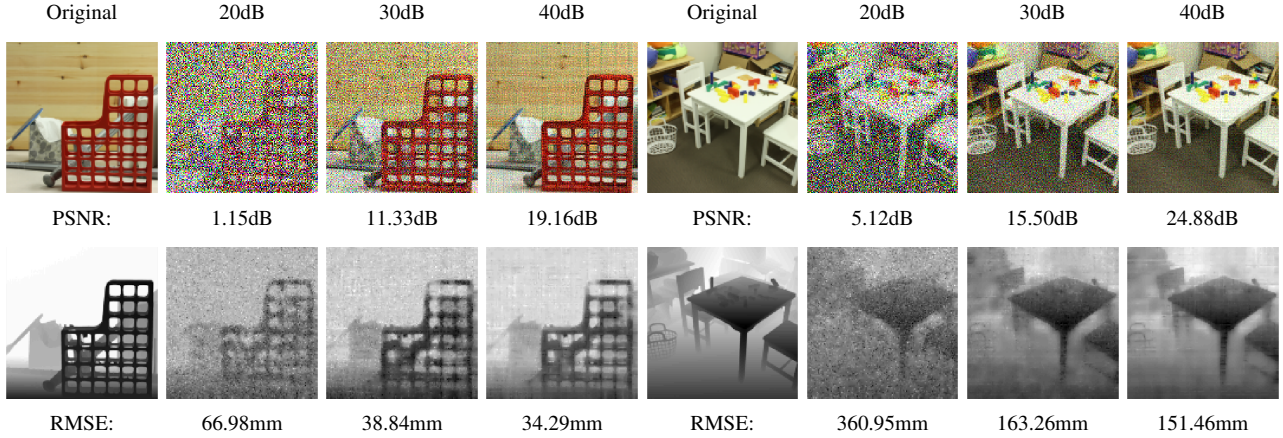


Fig. 5: Effects of noise: Reconstruction from the measurements with signal-to-noise ratio (SNR) at 20dB, 30dB and 40dB, along with the **PSNR** of reconstructed image and **RMSE** of reconstructed depth map. As expected, the quality of reconstructed image and depth improves as the noise level is reduced. The sequence in left is for *Sword*, right is *Playtable*.

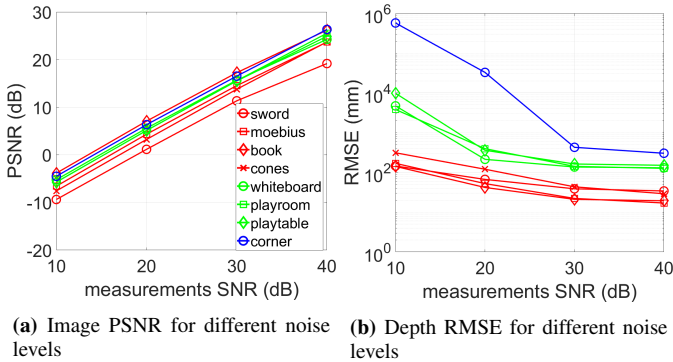


Fig. 6: Reconstruction from measurements with different levels of Gaussian noise on multiple scenes. Both of the image Peak Signal-to-Noise Ratio and depth Root mean squared error are improved as the noise is reduced. The reconstruction quality degrades if the scene is placed farther from the camera.

and all the points in the 3D volume can contribute to the sensor measurements without occluding one another. In contrast, we consider natural photographic scenes, where objects are usually opaque and block light from objects behind them along the same angular direction. We can model such scenes as having only one voxel along any angle to be nonzero; however, that will be a nonconvex constraint and to enforce that we will have to resort to some heuristic similar to the one in [6]. For the sake of comparison, we solve the ℓ_1 norm-based sparse recovery problem as described in [12], but then we pick the points with the maximum light intensity at each angle to form the reconstructed image and (inverse) depth map.

A comparison of different recovery methods with the same imaging setup is shown in Figure 8. For the same scene, we reconstruct the same measurements using the three methods. As we can observe that our proposed algorithm offers a significant improvement compared to existing methods in all the test scenes.

The time and storage complexity of our proposed method

and the other two methods depend on different factors; such as whether the imaging model is separable or convolutional and the sampling density along the depth. Since the main computational complexity of all the methods arises from the applications of the forward and adjoint operators, we will just discuss the complexity of those operators for different methods. The imaging operator in the greedy algorithm uses a separable mask and assumes that the scene consists of D depth planes. The computational complexity of the operator is $O(DMN^2)$ when we have $M \times M$ sensor pixels to reconstruct $N \times N$ image at D predefined depth planes. The convolutional model can be implemented using a fast Fourier transform and its complexity for a 3D volume with D depth planes is $O(DN^2 \log(N))$. The time complexity of forward imaging operator in the proposed method is $O(M^2N^2)$ because we assign independent depth values to each of the angles.

V. EXPERIMENTAL RESULTS

To demonstrate the performance of our proposed method in the real world, we built a FlatCam prototype to capture images of different objects with different depth profiles. Below we discuss the details of our experiments and present reconstructed intensity and depth maps for some real objects.

A. Prototype Setup

Image sensor. We used a Sony IMX249 CMOS color sensor that came inside a point grey camera (model BFLY-U3-23S6C-C). The sensor has 1920×1200 pixels and the size of each pixel is $5.86\mu\text{m}$. The physical size of the sensor is approximately $11.2\text{mm} \times 7\text{mm}$.

Mask pattern. We printed a binary mask pattern on a plastic sheet. The mask pattern was created by computing an outer product of two 255-length MLS vectors and setting all the -1 entries to 0. The physical size of each mask feature is $60\mu\text{m}$. The physical size of the generated mask pattern is approximately $15.3\text{mm} \times 15.3\text{mm}$.

Sensor and mask placement. We placed the mask and the bare sensor on two optical posts such that the mask-to-sensor

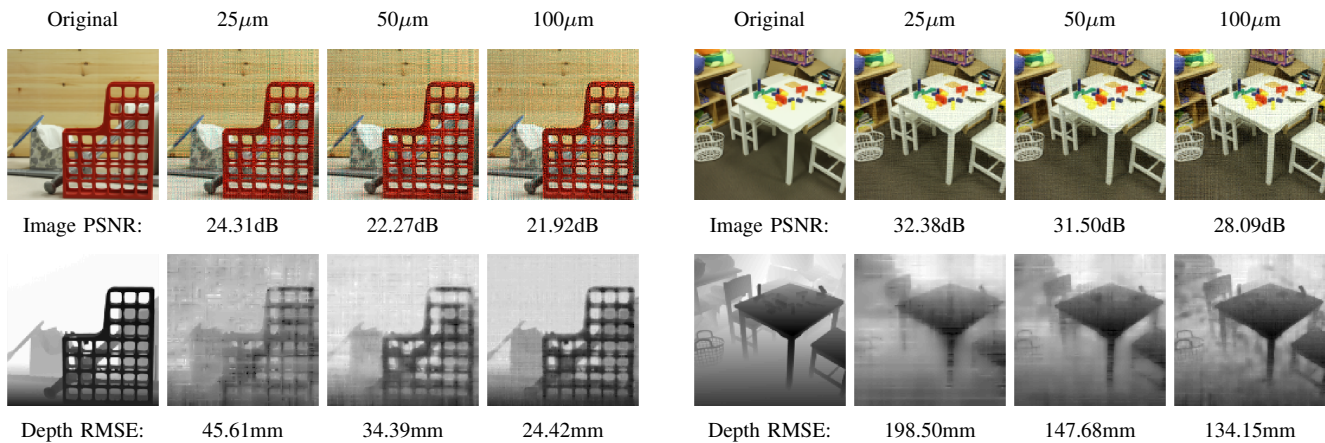


Fig. 7: Reconstructions from measurements with different sizes of sensor pixels. The number of sensor pixels is fixed as 512×512 . The quality of depth reconstruction improves as we increase the size of sensor pixels.

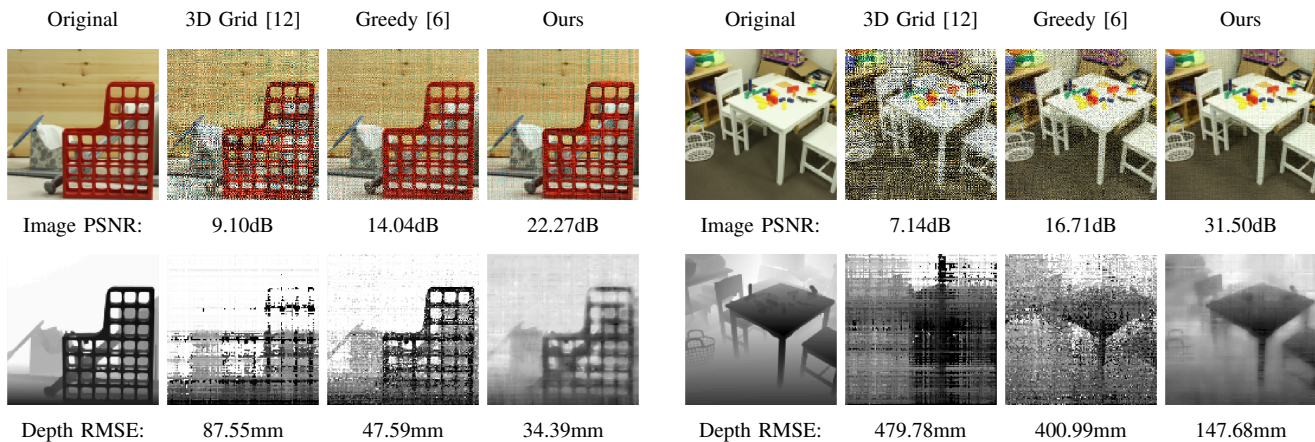


Fig. 8: Comparison of existing 3D recovery methods for lensless imaging, 3D grid method in [12], [17] and greedy method in [6], with our proposed method. 3D grid method provides a 3D volume with multiple depth planes; therefore, we pick the depth with the largest light intensity along any angle for comparison.

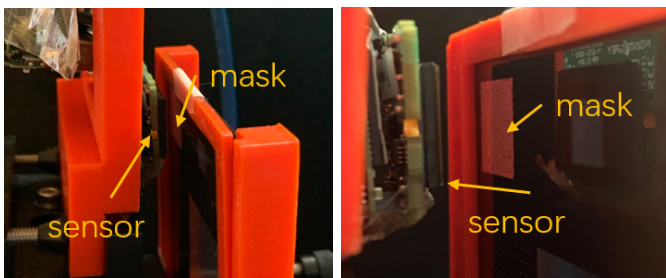


Fig. 9: Camera prototype. The side view of the sensor and mask assembly. The sensor and mask are placed at a large distance for this image, but their distance (d) is approximately 4mm in our experiments. The mask pattern is binary and separable, and the physical size of each feature is $60\mu\text{m}$.

distance (d) is approximately 4mm; we attached kinematic platforms on top of the optical posts so that we can align the sensor and mask. Pictures of our sensor and mask setup are shown in Figure 9.

Data acquisition and processing. In our experiments, we calibrated the system by capturing sensor measurements while moving an LED flashlight at different locations in front of the camera. We performed all our experiments by uniformly illuminating the object with a table lamp. We reconstructed depth map and colored images at 128×128 pixel resolution from 512×512 sensor measurements. The sensor provides 1920×1200 pixels; we first resize the sensor measurement into 960×600 pixels by binning 2×2 pairs, and then we crop a 512×512 area in the center.

B. Calibration of the Prototype Camera

We use a separable mask pattern and align the mask and sensor assembly such that the response of any point source on the sensor is a rank-one image after mean subtraction [5]. To calibrate system matrix at one given depth, we can capture a sequence of rank-1 Hadamard patterns as described in [5] or capture the response of one LED flashlight as described in [12]. Instead of calibrating the separable system matrices for different depth planes, we calibrated the mask pattern

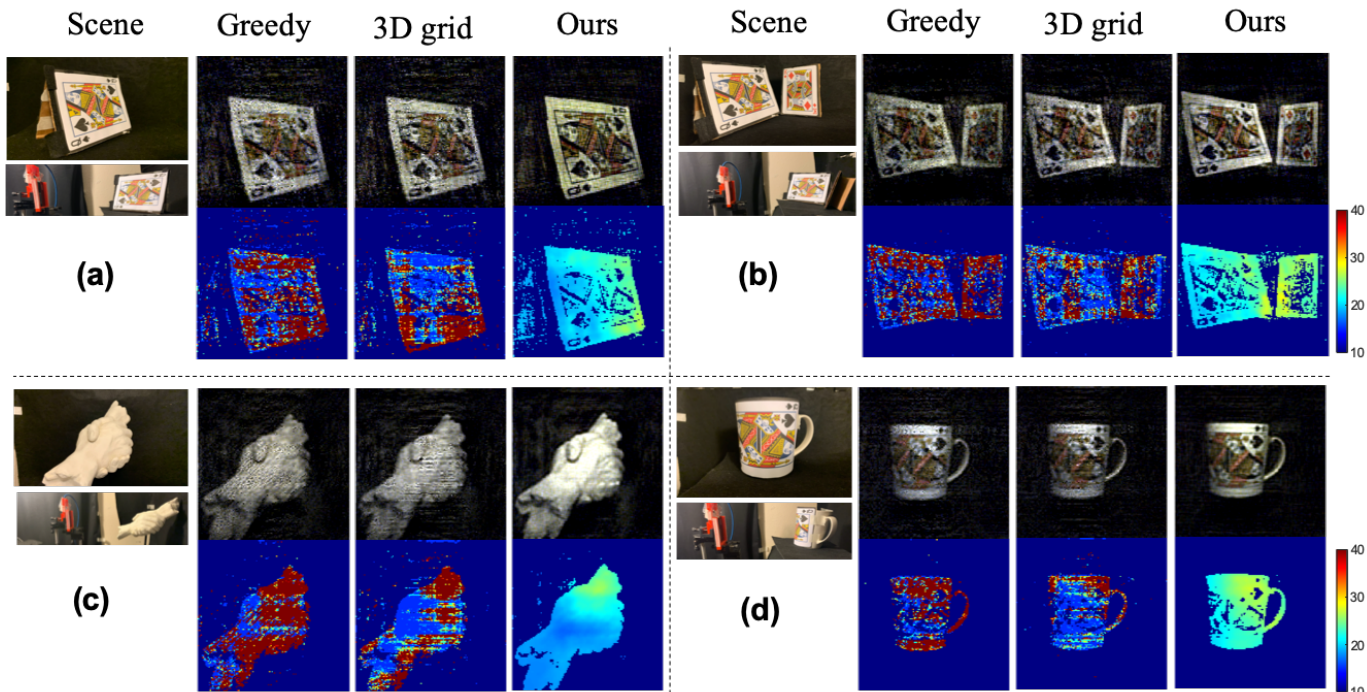


Fig. 10: Experiments on real objects. (a) A slanted card; the depth range is 18–28cm (b) Two slanted cards; the depth range of left card is 18–28cm and the right card is 26–29cm. (c) Hand sculpture; depth range is 15–30cm. (d) A mug with card texture; depth range is 24–27cm. We divide each group of real scenes into four columns, the first column is front view and side view of the scene, the second column is the result from greedy algorithm in [6], the third column is the output of sparse 3D grid recovery algorithm proposed in [12] and [17], and the last column is the image intensity and depth map estimated using our proposed algorithm.

function at one depth and evaluated the point spread function at arbitrary depth and angle according to (1). Because our mask pattern is bigger than the sensor, we captured sensor measurements for LED flashlight at 9 different angles at the same depth and merged them to estimate the mask function at that depth.

In our experiments, we captured the sensor measurements by placing an LED at $z = 42\text{cm}$ away from the sensor, which corresponds to the mask function in (1) evaluated at $\alpha = 1 - d/z = 0.9905$ for $d = 4\text{mm}$, $z = 42\text{cm}$. We first resized the calibrated mask function to compute the mask function corresponding to $\alpha = 1$.

C. Reconstruction of Real Objects

We present results for four objects in Figure 10, (a) slanted card has depth range from 18cm to 28cm, (b) two slanted cards have depth ranges from 18cm to 28cm and 26cm to 29cm, (c) hands sculpture has depth range from 15cm to 30cm, and (d) mug with card texture depth is from 24cm to 27cm. The figure is divided into four boxes. In each box, we present a front- and side-view of the object along with estimated scene intensity and depth maps for three methods: the greedy algorithm in [6], the sparse 3D volume recovery method from [12], [17], and our proposed method. For the greedy and 3D grid method, we generated 15 candidate depth planes by uniformly sampling the inverse depth values α between 0.96

and 0.9905 (corresponding to the depth of 10cm and 42cm, respectively).

All the objects in our experiments are placed in front of the black background and the depth values for dark pixels are not meaningful. We can observe that in all these experiments, our proposed method provides a continuous depth map that is consistent with the real depth of the object in the scene. In comparison, both the greedy algorithm [6] and the sparse 3D volume recovery algorithm [12], [17] produce coarse and discretized depth maps. The intensity map recovered by our method is also visually better compared to other methods.

Even though our proposed algorithm produces better intensity and depth maps compared to the greedy and 3D grid methods, we observed that the estimated depth has some errors in the darker parts of the objects. For instance, the left part of the mug is darker than the right part because the object was illuminated from a lamp on the right side. The left part appears to have errors in the depth estimate as several pixels are assigned small depth values, but that part is in fact farther from the sensor. We also observe a similar effect in other experiments, where depth estimates for darker parts of the scene appear to have larger errors.

VI. CONCLUSION

We presented a new algorithm to jointly estimate the image and depth of a scene using a single snapshot of a mask-based lensless camera. Existing methods for 3D lensless imaging

either estimate scene over a predefined 3D grid (which is computationally expensive) or a small number of candidate depth planes (which provides a coarse depth map). We divide the scene into an intensity map at uniform angles and a depth map on a continuous domain, which allows us to estimate a variety of scenes with different depth ranges using the same formulation. We jointly estimate the image intensity and depth map by solving a nonconvex problem. We initialize our estimates using a greedy method and add weighted regularization to enforce smoothness in the depth estimate while preserving the sharp edges. We demonstrated with extensive simulations and experiments with real data that our proposed method can recover image and depth with high accuracy for a variety of scenes. We evaluated the performance of our methods under different noise levels, sensor sizes, and numbers of sensor pixels and found the method to be robust. We presented a comparison with existing methods for lensless 3D imaging and demonstrated both in simulation and real experiments that our method provides significantly better results. We believe this work provides a step toward capturing complex scenes with lensless cameras, where depth estimation is a feature as well as a compulsion because if the depth information is unavailable or inaccurate, that will cause artifacts in the recovered images.

REFERENCES

- [1] Y. Zheng and M. S. Asif, "Image and depth estimation with mask-based lensless cameras," in *IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2019, pp. 91–95.
- [2] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [3] S. B. Gokturk, H. Yalcin, and C. Bamji, "A time-of-flight depth sensor - system description, issues and solutions," in *Conference on Computer Vision and Pattern Recognition Workshop*, June 2004, pp. 35–35.
- [4] F. Heide, M. B. Hullin, J. Gregson, and W. Heidrich, "Low-budget transient imaging using photonic mixer devices," *ACM Transactions on Graphics (ToG)*, vol. 32, no. 4, p. 45, 2013.
- [5] M. S. Asif, A. Ayremlou, A. Sankaranarayanan, A. Veeraraghavan, and R. G. Baraniuk, "Flatcam: Thin, lensless cameras using coded aperture and computation," *IEEE Transactions on Computational Imaging*, vol. 3, no. 3, pp. 384–397, Sept 2017.
- [6] M. S. Asif, "Lensless 3d imaging using mask-based cameras," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, April 2018, pp. 6498–6502.
- [7] A. Yedidia, C. Thrampoulidis, and G. Wornell, "Analysis and optimization of aperture design in computational imaging," *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 4029–4033, April 2018. [Online]. Available: <http://sigport.org/3049>
- [8] E. E. Fenimore and T. M. Cannon, "Coded aperture imaging with uniformly redundant arrays," *Appl. Opt.*, vol. 17, no. 3, pp. 337–347, Feb 1978. [Online]. Available: <http://ao.osa.org/abstract.cfm?URI=ao-17-3-337>
- [9] A. Busboom, H. Elders-Boll, and H. D. Schotten, "Uniformly redundant arrays," *Experimental Astronomy*, vol. 8, no. 2, pp. 97–123, Jun 1998. [Online]. Available: <https://doi.org/10.1023/A:1007966830741>
- [10] T. M. Cannon and E. E. Fenimore, "Coded Aperture Imaging: Many Holes Make Light Work," *Optical Engineering*, vol. 19, p. 283, Jun. 1980.
- [11] V. Boominathan, J. K. Adams, M. S. Asif, B. W. Avants, J. T. Robinson, R. G. Baraniuk, A. C. Sankaranarayanan, and A. Veeraraghavan, "Lensless imaging: A computational renaissance," *IEEE Signal Processing Magazine*, vol. 33, no. 5, pp. 23–35, 2016.
- [12] N. Antipa, G. Kuo, R. Heckel, B. Mildenhall, E. Bostan, R. Ng, and L. Waller, "Diffusercam: lensless single-exposure 3d imaging," *Optica*, vol. 5, no. 1, pp. 1–9, Jan 2018. [Online]. Available: <http://www.osapublishing.org/optica/abstract.cfm?URI=optica-5-1-1>
- [13] S. S. Khan, A. V. R., V. Boominathan, J. Tan, A. Veeraraghavan, and K. Mitra, "Towards photorealistic reconstruction of highly multiplexed lensless images," in *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [14] K. Monakhova, J. Yurtsever, G. Kuo, N. Antipa, K. Yanny, and L. Waller, "Learned reconstructions for practical mask-based lensless imaging," *Opt. Express*, vol. 27, no. 20, pp. 28 075–28 090, Sep 2019. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-27-20-28075>
- [15] A. Dave, A. K. Vadathya, R. Subramanyam, R. Baburajan, and K. Mitra, "Solving inverse computational imaging problems using deep pixel-level prior," *IEEE Trans. Computational Imaging*, vol. 5, no. 1, pp. 37–51, 2019. [Online]. Available: <https://doi.org/10.1109/TCL.2018.2882698>
- [16] H. H. Barrett, D. T. Wilson, G. D. DeMeester, and H. Scharfman, "Fresnel zone plate imaging in radiology and nuclear medicine," *Optical Engineering*, vol. 12, no. 1, pp. 8 – 12 – 5, 1973. [Online]. Available: <https://doi.org/10.1117/12.7971622>
- [17] 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," *Science Advances*, vol. 3, no. 12, 2017. [Online]. Available: <https://advances.sciencemag.org/content/3/12/e1701548>
- [18] Y. Hua, S. Nakamura, M. S. Asif, and A. C. Sankaranarayanan, "Sweepcam depth-aware lensless imaging using programmable masks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 7, pp. 1606–1617, 2020.
- [19] Z. Tan, P. Yang, and A. Nehorai, "Joint sparse recovery method for compressed sensing with structured dictionary mismatches," *IEEE Transactions on Signal Processing*, vol. 62, no. 19, pp. 4997–5008, Oct 2014.
- [20] N. Boyd, G. Schiebinger, and B. Recht, "The alternating descent conditional gradient method for sparse inverse problems," in *IEEE 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, Dec 2015, pp. 57–60.
- [21] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Transactions on Information Theory*, vol. 53, no. 12, pp. 4655–4666, Dec 2007.
- [22] D. Needell and J. A. Tropp, "Cosamp: Iterative signal recovery from incomplete and inaccurate samples," *Commun. ACM*, vol. 53, no. 12, pp. 93–100, Dec. 2010. [Online]. Available: <http://doi.acm.org/10.1145/1859204.1859229>
- [23] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1982–2001, 2010.
- [24] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The convex geometry of linear inverse problems," *Foundations of Computational Mathematics*, vol. 12, no. 6, pp. 805–849, Dec 2012. [Online]. Available: <https://doi.org/10.1007/s10208-012-9135-7>
- [25] G. Tang, B. N. Bhaskar, P. Shah, and B. Recht, "Compressed sensing off the grid," *IEEE Transactions on Information Theory*, vol. 59, no. 11, pp. 7465–7490, Nov 2013.
- [26] Z. Yang, L. Xie, and C. Zhang, "Off-grid direction of arrival estimation using sparse bayesian inference," *IEEE Transactions on Signal Processing*, vol. 61, no. 1, pp. 38–43, Jan 2013.
- [27] D. Takhar, J. N. Laska, M. B. Wakin, M. F. Duarte, D. Baron, S. Sarvotham, K. F. Kelly, and R. G. Baraniuk, "A new compressive imaging camera architecture using optical-domain compression," *Proc.SPIE*, vol. 6065, pp. 6065 – 6065 – 10, 2006. [Online]. Available: <https://doi.org/10.1117/12.659602>
- [28] A. Zomet and S. K. Nayar, "Lensless imaging with a controllable aperture," in *IEEE Computer Vision and Pattern Recognition*, vol. 1, June 2006, pp. 339–346.
- [29] Y. Zheng and M. S. Asif, "Imaging with distributed lensless line sensors," in *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, Nov 2019, pp. 1289–1293.
- [30] 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 Processing Magazine*, vol. 25, no. 2, pp. 83–91, March 2008.
- [31] D. Reddy, J. Bai, and R. Ramamoorthi, "External mask based depth and light field camera," in *2013 IEEE International Conference on Computer Vision Workshops*, Dec 2013, pp. 37–44.
- [32] A. Levin, R. Fergus, F. Durand, and W. T. Freeman, "Image and depth from a conventional camera with a coded aperture," *ACM Trans. Graph.*, vol. 26, no. 3, Jul. 2007. [Online]. Available: <http://doi.acm.org/10.1145/1276377.1276464>

- [33] A. Veeraraghavan, R. Raskar, A. Agrawal, A. Mohan, and J. Tumblin, "Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing," in *ACM transactions on graphics (TOG)*, vol. 26, no. 3. ACM, 2007, p. 69.
- [34] K. Marwah, G. Wetzstein, Y. Bando, and R. Raskar, "Compressive Light Field Photography using Overcomplete Dictionaries and Optimized Projections," *ACM Trans. Graph. (Proc. SIGGRAPH)*, vol. 32, no. 4, pp. 1–11, 2013.
- [35] M. Hirsch, S. Sivaramakrishnan, S. Jayasuriya, A. Wang, A. Molnar, R. Raskar, and G. Wetzstein, "A switchable light field camera architecture with angle sensitive pixels and dictionary-based sparse coding," in *2014 IEEE International Conference on Computational Photography (ICCP)*, 2014, pp. 1–10.
- [36] Y. Wu, V. Boominathan, H. Chen, A. Sankaranarayanan, and A. Veeraraghavan, "Phasecam3d learning phase masks for passive single view depth estimation," in *2019 IEEE International Conference on Computational Photography (ICCP)*, May 2019, pp. 1–12.
- [37] J. Chang and G. Wetzstein, "Deep optics for monocular depth estimation and 3d object detection," in *Proc. IEEE ICCV*, 2019.
- [38] X. Lin, Y. Rivenson, N. T. Yardimci, M. Veli, Y. Luo, M. Jarrahi, and A. Ozcan, "All-optical machine learning using diffractive deep neural networks," *Science*, vol. 361, no. 6406, pp. 1004–1008, 2018. [Online]. Available: <https://science.sciencemag.org/content/361/6406/1004>
- [39] D. Mengu, Y. Luo, Y. Rivenson, and A. Ozcan, "Analysis of diffractive optical neural networks and their integration with electronic neural networks," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 26, no. 1, pp. 1–14, Jan 2020.
- [40] M. Schmidt, "minfunc: unconstrained differentiable multivariate optimization in matlab," 2005. [Online]. Available: <http://www.cs.ubc.ca/~schmidtm/Software/minFunc.html>
- [41] D. C. Liu and J. Nocedal, "On the limited memory bfgs method for large scale optimization," *Mathematical Programming*, vol. 45, no. 1, pp. 503–528, Aug 1989. [Online]. Available: <https://doi.org/10.1007/BF01589116>
- [42] Y. Liu, J. Ma, Y. Fan, and Z. Liang, "Adaptive-weighted total variation minimization for sparse data toward low-dose x-ray computed tomography image reconstruction," *Physics in Medicine & Biology*, vol. 57, no. 23, p. 7923, 2012. [Online]. Available: <http://stacks.iop.org/0031-9155/57/i=23/a=7923>
- [43] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)*, Jan 1998, pp. 839–846.
- [44] F. Durand and J. Dorsey, "Fast bilateral filtering for the display of high-dynamic-range images," *ACM Trans. Graph.*, vol. 21, no. 3, pp. 257–266, Jul. 2002. [Online]. Available: <http://doi.acm.org/10.1145/566654.566574>
- [45] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D*, vol. 60, no. 1-4, pp. 259–268, Nov. 1992. [Online]. Available: [https://doi.org/10.1016/0167-2789\(92\)90242-F](https://doi.org/10.1016/0167-2789(92)90242-F)
- [46] T. Goldstein and S. Osher, "The split bregman method for l_1 -regularized problems," *SIAM J. Img. Sci.*, vol. 2, no. 2, pp. 323–343, Apr. 2009. [Online]. Available: <http://dx.doi.org/10.1137/080725891>
- [47] F. J. MacWilliams and N. J. A. Sloane, "Pseudo-random sequences and arrays," *Proceedings of the IEEE*, vol. 64, no. 12, pp. 1715–1729, Dec 1976.
- [48] D. Scharstein, R. Szeliski, and R. Zabih, "A taxonomy and evaluation of dense two-frame stereo correspondence algorithms," in *Proceedings IEEE Workshop on Stereo and Multi-Baseline Vision*, Dec 2001, pp. 131–140.
- [49] A. Saxena, S. H. Chung, and A. Y. Ng, "Learning depth from single monocular images," in *Advances in Neural Information Processing Systems 18*, Y. Weiss, B. Schölkopf, and J. C. Platt, Eds. MIT Press, 2006, pp. 1161–1168. [Online]. Available: <http://papers.nips.cc/paper/2921-learning-depth-from-single-monocular-images.pdf>
- [50] A. Saxena, M. Sun, and A. Y. Ng, "Make3d: Learning 3d scene structure from a single still image," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 5, pp. 824–840, May 2009.
- [51] P. K. Nathan Silberman, Derek Hoiem and R. Fergus, "Indoor segmentation and support inference from rgb-d images," in *ECCV*, 2012.



Yucheng Zheng received his B.Sc. degree in electrical engineering from the Nanjing University of Aeronautics and Astronautics, Nanjing, China in 2017. He is currently pursuing a Ph.D. degree at the University of California, Riverside, CA, USA. His research interests include computational imaging, computer vision and signal processing.



M. Salman Asif received the B.Sc. degree in 2004 from the University of Engineering and Technology, Lahore, Pakistan, and the M.S.E.E degree in 2008, and the Ph.D. degree in 2013 from the Georgia Institute of Technology, Atlanta, GA, USA. He is an Assistant Professor in the Department of Electrical and Computer Engineering, University of California, Riverside, CA, USA. He was a Research Intern at Mitsubishi Electric Research Laboratories, Cambridge, MA, USA, in the Summer of 2009, and at Samsung Standards Research Laboratory, Richardson, TX, USA, in the Summer of 2010. He was a Senior Research Engineer at Samsung Research America, Dallas, TX, USA, from August 2012 to January 2014, and as a Postdoctoral Researcher at Rice University from February 2014 to June 2016. His research interests include compressive sensing, computational and medical imaging, and machine learning.