



End-to-end learning for joint depth and image reconstruction from diffracted rotation

Mazen Mel¹ · Muhammad Siddiqui² · Pietro Zanuttigh¹

Accepted: 20 October 2023
© The Author(s) 2023

Abstract

Monocular depth estimation is an open challenge due to the ill-posed nature of the problem at hand. Deep learning techniques proved capable of producing acceptable depth estimation accuracy but the lack of robust depth cues within RGB images severally limits their performance. Coded aperture-based methods using phase and amplitude masks encode strong depth cues within 2D images by means of depth-dependent Point Spread Functions (PSFs) at the price of a reduced image quality. In this paper, we propose a novel end-to-end learning approach for depth from diffracted rotation. A phase mask that produces a Rotating Point Spread Function (RPSF) as a function of defocus is jointly optimized with the weights of a depth estimation neural network. To this aim, we introduce a differentiable physical model of the aperture mask and exploit an accurate simulation of the camera imaging pipeline. Our approach requires a significantly less complex model and less training data, yet it outperforms existing methods for monocular depth estimation on indoor benchmarks. In addition, we address the image degradation problem by incorporating a non-blind and nonuniform image deblurring module to recover the sharp all-in-focus image from its blurred counterpart.

Keywords Monocular depth estimation · RPSFs · Image deblurring

1 Introduction

Depth estimation from a single RGB image is an ill-posed inverse problem, thus, additional image priors or sophisticated imaging systems are generally needed to account for the lack of depth cues in captured images.

For example, the camera optics may be modified accordingly, a more advanced image sensor design and task-specific post processing algorithms may be developed in order to capture information well beyond the capabilities of conventional imaging systems.

In this paper, we propose a computational camera wherein the Point Spread Function (PSF) is altered in order to encode depth information within a 2D image. The desired PSF is obtained via an optimized phase mask inserted at the aperture plane of the camera. Our system encodes meaningful

and robust depth cues in single RGB images thus making it easier for post-processing algorithms to produce accurate depth data. However, such approach generally suffers from image quality degradation due to the poor light efficiency and/or the low Modulation Transfer Function (MTF) of the engineered PSF for high spatial frequency components. In fact, amplitude aperture masks block a significant amount of light from reaching the image sensor resulting in low light throughput, while using phase only masks solves such problem since it only acts on the phase component of the incoming light wave. Still, the MTF drops rapidly in high frequency regions and the depth-dependent blurring caused by the camera's engineered PSF produces low SNR and degrade image quality.

We propose a full pipeline that takes as input a single RGB image and produces an estimated depth map of the scene along with the recovered sharp image from its PSF-blurred counterpart. We introduce a novel end-to-end deep learning model that jointly deals with the PSF design optimization and the depth estimation task. To this aim, a differentiable physical model of the aperture mask is introduced together with an accurate simulation of the imaging pipeline including the optimized optics. In this way the learned model is

✉ Mazen Mel
mazen.mel@phd.unipd.it

¹ Department of Information Engineering, University of Padova, Padova, Italy

² Sony Europe B.V., Zweigniederlassung Deutschland Stuttgart Technology Center, Stuttgart, Germany

able to firstly predict the optimal parameters of the phase mask design and then estimate the depth data from the RPSF-blurred input. Finally, in order to address the image quality issue, a deep learning based non-blind and nonuniform deblurring module is incorporated.

2 Related work

Rotating Point Spread Functions RPSFs have been theoretically shown to increase the Fisher Information along the depth dimension by at least an order of magnitude as they have a uniformly lower Cramér-Rao bound across the axial dimension [1], which makes them highly sensitive to depth changes. RPSFs are obtained using pupil engineered cameras by means of phase and/or amplitude masks that are inspired by the concept of Orbital Angular Momentum (OAM) of light beams [2]. A beam with a rotating light intensity distribution along the propagation axis can be generated by a linear superposition of a set of suitably chosen Gaussian-Laguerre (GL) modes [1, 3] that can be optically encoded by the aperture mask. These masks generate a PSF with invariant features that continuously rotate with defocus. However, such approach suffers from poor light throughput, some works [4, 5] addressed this problem and proposed iterative optimization schemes to ensure higher light efficiency of the engineered RPSF. Depth dependent RPSFs can be also generated by phase only masks. For instance, a phase mask design inspired by Spiral Phase Plates (SPPs) [6, 7] was introduced in [8] where the pupil is subdivided into a set of annular Fresnel zones with an azimuthally increasing thickness profile, the delay imposed on the incident light waves increases azimuthally generating a corkscrew like wave-front carrying an OAM with a rotating phase function. In [9] and [10], the authors generalized the previous phase mask design by considering a phase function that allows for generating multi-order-helix RPSFs by introducing new design parameters: the number of rotating lobes within the RPSF and the confinement of each zone, i.e., the inner and outer radii of each annular region, in addition to the number of Fresnel zones. While [9] and [10] used purely empirical approaches to determine the values for each design parameter, in this work we optimize those parameters in an end-to-end fashion jointly with the weights of a depth estimation deep neural network.

Depth estimation using coded apertures Levin et al. [11] proposed an amplitude modulation mask that was placed in front of the camera lens to encode depth via the diffracted pattern of the camera's PSF: as point-like sources move along a plane parallel to that of the camera sensor, the mask shadow would shift accordingly and as they move closer or farther from the camera, the pattern would expand or shrink. This information is later used to determine the object's distance from the camera. The mask introduced by [11] has opaque

regions blocking a significant amount of light from reaching the image sensor. Zhou et al. [12] built upon [11] and introduced two complementary amplitude masks.

In [1] the concept of depth from diffracted rotation was introduced: a superposition of a set of suitably chosen Gaussian-Laguerre modes generates a double helix RPSF that was used to estimate the depth of a planar scene by analyzing the blurring effects within the captured image. However, low MTF by the mask leads to poor SNR within the captured image, thus limiting the capability of signal-processing based algorithms to recover sharp images or accurate depth maps. In addition, earlier studies relied on a design paradigm based on the separate optimization of the camera optics and post-processing algorithms: they design the mask first and then tailor a reconstruction algorithm that fits the proposed physical design of the mask as in [10, 13–15]. Such design methodology, however, leads to sub-optimal performance.

End-to-end learning for monocular depth estimation Recently, an emerging trend appeared to tackle the separate design issue and new frameworks for joint optical and digital optimization using deep learning techniques have been introduced. These methods exploit end-to-end learning to optimize the mask's height map together with the trainable weights of a Convolutional Neural Network (CNN). Haim et al. [16] proposed a differentiable phase mask consisting of concentric rings that introduce depth-dependent chromatic aberrations and encoding depth cues within single captured images. In the work of Chang et al. [17] and Wu et al. [18] (that exploits a layered depth map to generate the coded image) a free-form differentiable phase mask design parameterized using a set of superposed Zernike polynomials [19, 20] was jointly optimized with the weights of a U-Net [21]. However, the employed camera model was not realistic accounting only for additive Gaussian noise [18]. Furthermore, unrestricted parameterized mask design and higher degrees of freedom may cause the optimization to converge toward local minimas as the objective function becomes too complex leading to sub-optimal performance. In this work, the mask is parameterized using only three design parameters two of which are optimized in an end-to-end fashion, and the PSF has clear and simple correlations with depth.

Image deblurring Non-blind image deblurring has been extensively studied before (e.g., [22–25]), and a substantial performance gain has been made easier using deep learning based image deblurring techniques.

The authors of [26] used the separability property of the pseudo-inverse kernel in Wiener deconvolution filter to design a dedicated CNN. In [27], a two-stage deblurring process was introduced using a Wiener deconvolution filter and a simple MLP architecture. A novel deconvolution approach was recently introduced by [28] where a Wiener deconvolution filter is applied to the input data in feature space.

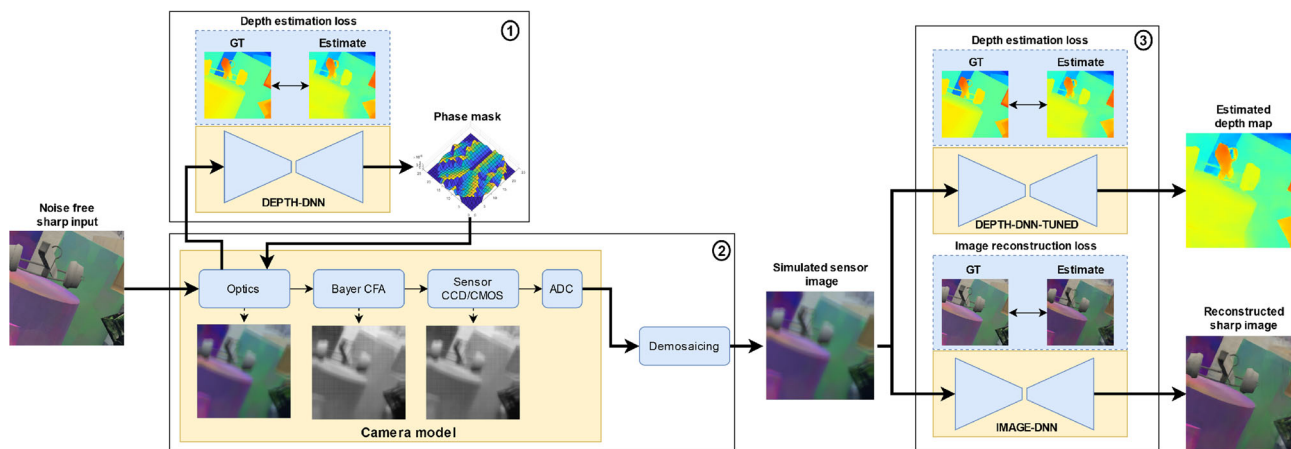


Fig. 1 The full architecture of our end-to-end learning framework

3 Depth estimation from diffracted rotation

We propose an end-to-end learning approach for depth from diffracted rotation using RPSF-coded images. As shown in Fig. 1, the full pipeline of the proposed solution encompasses three main stages. In the first stage, the height map of the phase mask is jointly learned with the weights of a neural network (DEPTH-DNN) trained to perform monocular depth estimation. This module is trained using noise-free RPSF-blurred synthetic images. In the second stage, the optimized phase mask is fitted within the optics module and a digital image formation pipeline is applied to the RPSF-coded synthetic images in order to simulate a realistic camera model. Finally, we used the demosaiced images from the previous stage as input to fine-tune the weights of DEPTH-DNN on noisy data obtaining a refined model (DEPTH-DNN-TUNED) and to recover the all-in-focus sharp image using a dedicated network (IMAGE-DNN) which performs non-blind and nonuniform image deblurring. Both of these modules make up the third and last stage of the proposed architecture.

3.1 Engineered PSF

In this section, we briefly describe the effect of modified optics on the system’s PSF (more details are in the supplementary material). For simplicity, consider an on-axis ideal point source situated at optical infinity, the light field U_{in} with a constant amplitude P and a phase ψ emanating from such source has the form: $U_{in} = P e^{i\psi}$. An optical element (e.g., a lens or a phase mask) with a refractive index n and a height profile h introduces a phase delay $\Phi = \frac{2\pi(n-1)}{\lambda}h$ on incident light wave-fronts. If a phase mask is inserted at the entrance of an imaging system, the total phase delay can be

expressed as the sum of the delay due to the lens with the one due to the mask:

$$\Phi_{optics} = \Phi_{lens} + \Phi_{mask} \tag{1}$$

The light field after the lens and mask system has the form:

$$U_{out} = A \cdot P \cdot e^{i(\psi + \Phi_{optics})} \tag{2}$$

Where A is the aperture mask simulating the finite aperture area. The field U_{out} can be further propagated to the image plane and the PSF can be obtained by the field’s intensity distribution. Choosing the appropriate height profile of the phase mask helps designing specific PSF patterns depending on the target task.

3.2 Phase mask design

Both annular and free-form mask designs have been studied in the context of joint optimization of camera optics with post-processing algorithms [16–18]. We build upon the depth-dependent RPSF introduced by [8]. In this section, the mathematical model of the mask’s phase profile is described and in the next section a differentiable approximation of the mask’s height map is derived to allow for gradient back-propagation in our end-to-end learning framework.

We use a Fresnel-zone-based design [8] that has an outermost radius R and L phase plates in the form of concentric annular regions each of topological charge $l = 1, \dots, L$ and bounded by two radii $R_{l-1} = R\sqrt{\frac{l-1}{L}}$ and $R_l = R\sqrt{\frac{l}{L}}$.

The pupil plane phase can be written as in [8]:

$$\Phi_{mask}(\rho, \phi) = \begin{cases} \phi & 0 \leq \rho < \sqrt{\frac{1}{L}} \\ \vdots & \\ l\phi & \sqrt{\frac{l-1}{L}} \leq \rho < \sqrt{\frac{l}{L}} \\ \vdots & \\ L\phi & \sqrt{\frac{L-1}{L}} \leq \rho < 1 \end{cases} \quad (3)$$

Φ_{mask} is defined with the polar position vector ρ normalized by the pupil's outermost radius R , and is a step function modeling the physical design property of concentric rings each with its own phase profile. The phase profile in Eq. 3 can be generalized to account for multi-order-helix RPSFs as in [10] and [9] by expanding it into:

$$\Phi_{mask}(\rho, \phi) = \begin{cases} \phi & 0 \leq \rho < (\frac{1}{L})^\epsilon \\ \vdots & \\ [(l-1)N+1]\phi & (\frac{l-1}{L})^\epsilon \leq \rho < (\frac{l}{L})^\epsilon \\ \vdots & \\ [(L-1)N+1]\phi & (\frac{L-1}{L})^\epsilon \leq \rho < 1 \end{cases} \quad (4)$$

Notice that the inner and outermost radii of each zone are now controlled by ϵ which lies in $[0, 1]$, and the topological charge of each ring is now $[(l-1)N+1]$ instead of just l . Besides the number of rings L , N and ϵ are two new design parameters each having an effect over the resulting RPSF shape. More precisely, N defines the number of peaks or the main rotating lobes of the RPSF and L and ϵ both control the peak separation and confinement of each peak. In the case of a single helix rotating PSF ($N = 1$) the phase profile of each ring would be reduced to the original expression as in Eq. 3. Notice also that by increasing the number of peaks, the practical depth range would be reduced because of rotation ambiguities when peaks rotate beyond $[-\frac{\pi}{N}, \frac{\pi}{N}]$.

3.3 Differentiable phase mask design

The phase profile presented in Eq. 4 is not differentiable with respect to the design parameters N and ϵ . Thus, a differentiable approximation for this equation is necessary in order to simulate the camera's optical layer and enable both forward and backward propagation. The number of Fresnel zones L is considered as a hyper-parameter that can be manually tuned to achieve better depth estimation performance.

The steps in Eq. 4 can be approximated with a set of 2D rings in polar coordinates with increased radii as L increases as illustrated in Fig. 2, each Fresnel zone is obtained by subtracting the areas of two 2D tanh functions each with a radius

corresponding to the one of the two radii of the desired ring. We multiply the inner coordinates of each tanh by a large constant (100 in our case) in order to get sharp mask edges. The new approximated phase profile equation for L zones can be written as:

$$\tilde{\Phi}_{mask}(\rho, \phi) = \sum_{l=0}^{l=L} \underbrace{((l-1)N+1)\phi}_{l^{th} \text{ ring phase}} \times \underbrace{\frac{1}{2}(\tanh[100(\rho-r_l)] - \tanh[100(\rho-r_{l+1}]))}_{l^{th} \text{ ring mask}} \quad (5)$$

Where $r_0 = 0$, $r_l = R(\frac{l}{L})^\epsilon$; $\forall l \in [1..L]$, and R is the outermost radius of the mask. ϕ and ρ are the polar coordinates.

Each phase profile $\Phi_{mask}^l = ((l-1)N+1)\phi$ is multiplied by the corresponding ring mask and the resulting zones are added up to produce the final phase profile of the mask. We produce the height map h of the mask as follows:

$$h(x, y) = \frac{\lambda}{2\pi(n-1)} \cdot (\arg\{e^{i\tilde{\Phi}_{mask}}\} \bmod 2\pi) \quad (6)$$

where \arg is the complex argument function, and the modulo accounts for the phase wrapping operation (phase values have a 2π periodicity as shown in Fig. 3).

Notice that the height profile of the phase mask is wavelength dependent: the resulting PSFs for the three RGB color primaries have different rotation rates which introduces chromatic aberrations. Still, it will not be problematic in the context of an end-to-end optimization framework since the network could learn the correlations between the rotation rate and the corresponding color channel. In fact, such aberrations can be seen as depth-dependent and can also relay valuable depth cues. For the physical design of the mask a reference wavelength value ($\lambda = 536.67$ nm) is used to produce the height map.

Figure 4 shows a double helix RPSF generated by a phase mask with $[N = 2, L = 5, \epsilon = 0.9]$, it has two main lobes rotating counterclockwise as a function of defocus. Notice that even at the in-focus plane the RPSF has the same shape thus objects that are "in-focus" are also blurred.

4 Training data generation

4.1 Datasets

A subset of FlyingThings3D dataset is used for the joint optimization of the phase mask and depth estimation network, and to train IMAGE-DNN. This dataset is a part of the Scene Flow synthetic datasets [29]. This subset was previously used by [18] in a similar joint optimization approach,

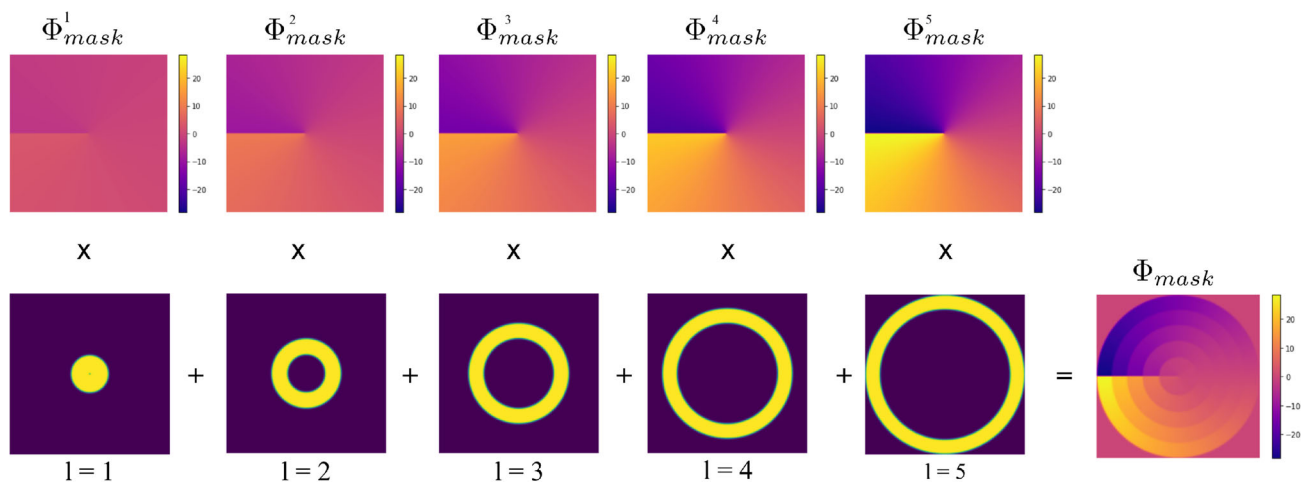


Fig. 2 A differentiable phase mask design where each Fresnel zone is simulated by a ring mask multiplied by the phase profile (Φ_{mask}^i) corresponding to each zone




Fig. 3 The height map of the phase mask. From left to right, the phase distribution as the argument of $e^{i\tilde{\psi}}$, the $[0, 2\pi]$ wrapped phase distribution, and the obtained height map for $[N = 2, L = 5, \epsilon = 0.9]$

it contains synthetic images of randomly placed objects with pixel-accurate disparity maps. The training, validation, and test sets contain respectively 5078, 555, and 420 images with a resolution of 278×278 pixels. Additionally, in order to evaluate the performance of our approach with the state-of-the-art in the task of monocular depth estimation on real data, NYUV2 [30] depth dataset is used to train and evaluate DEPTH-DNN. Originally, the dataset contains 120k training images of indoor scenes with a resolution of 640×480 pixels along with ground truth depth maps acquired by a Microsoft Kinect V1 depth sensor. The test set, as defined by the split in [31], contains 654 images. In this work, a subset of 50k samples of NYUV2 is used to train the depth estimation network as in [32]. Finally, the test set of SUNRGBD dataset [33] is used to evaluate the generalization capability of DEPTH-DNN. This set contains 5050 test images of indoor scenes with ground truth depth maps acquired by four different depth sensors some of which use active illumination techniques and others incorporate passive stereo systems.

4.2 Image formation model

Light rays emanating from the scene are acquired by the camera and optically coded by the phase mask via a depth-dependent blurring process with the camera’s RPSF. For our simulations, the ground truth depth maps are approximated with a layered model in which only a finite number of depth planes are used to compose and render the RPSF-coded image using the following image formation model:

$$I_{sim} = \sum_{d=1}^{d=D} (I_{aif} * RPSF_d) \odot M_d \tag{7}$$

Where I_{sim} is the simulated blurred image, I_{aif} is its all-in-focus counterpart, $RPSF_d$ is the RPSF intensity distribution at the depth plane d , \odot stands for element-wise multiplication, and $\{M_d; d = 1, \dots, D\}$ are the depth masks presenting the individual depth layers such that at each pixel location




Fig. 4 A double helix RPSF obtained by a mask with $[N = 2, L = 5, \epsilon = 0.9]$. The RPSF's intensity distributions are shown as a function of defocus

$\sum_d M_d = 1$, i.e., only one pixel mask is set to one at each position.

Afterward, the image formation pipeline is applied to the RPSF-blurred images to simulate real cameras. As illustrated by Fig. 1, a Bayer CFA receives the full color resolution RPSF-blurred image and produces down-sampled RGB color pattern. Although it is hard to accurately simulate the noise behavior within the sensor chip, the final amount of noise is mainly caused by sensor shot and read noises. To this end, we simulated the read noise with an additive Gaussian $\mathcal{N}(0, \sigma^2)$ with zero mean and a fixed standard deviation $\sigma = 0.01$, photon shot noise follows a Poisson distribution, in practice it is modeled by a Gaussian distribution whose mean and variance depend on the expected photon count over the exposure time. The resulting noisy sensor image is quantized with an ADC module with a resolution of 8 bits. Finally, a linear interpolation-based demosaicing technique of [34] is used to recover the full color channels from the CFA and produce the final output which will be used to fine-tune DEPTH-DNN and train IMAGE-DNN.

5 Deep learning framework

5.1 Monocular depth estimation

In the first stage of the proposed solution (see Fig. 1), the phase mask design parameters $[N, \epsilon]$ are jointly learned with the weights of a U-Net [21] which is trained on a sub-

set of FlyingThings3D [29]. Two different learning rates, $L_r^{mask} = 0.1$ and $L_r^{dnn} = 1e - 4$ are set for the phase mask and for the depth estimation neural network. During the training process, the gradient error is back-propagated through the network as well as the mask's trainable parameters and the weights are updated accordingly using TensorFlow's automatic differentiation framework [35]. The network of this stage is trained for 150k iterations using a batch size of 20. Adam optimizer [36] is used with exponential decay rates of the first momentum and second momentum, respectively, set to $\beta_1 = 0.99$ and $\beta_2 = 0.999$. The training takes roughly 12h using a single NVIDIA TITAN X GPU.

Similar to Wu et al. [18], we used a combination of Root Mean Square Error (RMSE) loss \mathcal{L}_{rmse} and gradient loss \mathcal{L}_{grad} which forces the network to estimate accurate depth maps with well defined object boundaries at different depth planes. The total loss function is defined as:

$$\mathcal{L}_{depth} = \mathcal{L}_{rmse} + \mathcal{L}_{grad} \quad (8)$$

where \mathcal{L}_{rmse} and \mathcal{L}_{grad} are defined as:

$$\mathcal{L}_{rmse}(\theta, \theta^*) = \sqrt{\frac{1}{|T|} \sum_{\theta \in T} (\theta - \theta^*)^2} \quad (9)$$

$$\mathcal{L}_{grad}(\theta, \theta^*) = \mathcal{L}_{rmse} \left(\frac{\partial \theta}{\partial x}, \frac{\partial \theta^*}{\partial x} \right) + \mathcal{L}_{rmse} \left(\frac{\partial \theta}{\partial y}, \frac{\partial \theta^*}{\partial y} \right) \quad (10)$$

θ and θ^* are, respectively, the predicted and the ground truth disparity maps while the subtraction is done pixel-by-pixel, x and y are the spatial dimensions, and $|T|$ is the number of disparity maps.

In the third stage, the same U-Net [21] which was previously trained on noise-free RPSF data is fine-tuned using 50k training iterations with noisy images simulated by the camera model of the second stage (see Fig. 1), the phase mask is fixed during this training pass and only the network's weights are updated. All hyper-parameters' values are fixed to the same values as in the first training stage.

For NYUV2 dataset [30], both the network and the mask are learned using 50k training samples, the input images to the network and the output depth maps have a resolution of 320×240 which correspond to half of that of the original samples, a bilinear-upsampling is applied to the predicted depth maps to recover the original resolution of 640×480 for evaluation purposes as in [32, 37, 38]. The network is trained for 150k iterations with a batch size of 20, the phase mask and the neural network learning rates as well as the optimizer used in the training are the same as the ones used for the subset of FlyingThings3D [29].

5.2 Image deblurring

We used a modified version of the deblurring model proposed by [28] wherein a Wiener deconvolution is applied to a set of feature maps extracted from the blurred input using a simple CNN architecture as feature extractor, the deconvolved feature maps are then fed into a multi-scale feature refinement stage in order to get the final deblurred image in a coarse-to-fine based reconstruction technique. Such approach proved capable of restoring very fine structural details allowing for accurate image reconstruction. In this work, we adapted the network proposed by [28] for the case of nonuniform image deblurring as the blur kernel in our case is spatially variant. More precisely, a separate deconvolution is performed for every depth layer and the results are cropped using the corresponding depth masks M_d as in Eq. 7. The deconvolved feature maps are then combined in order to get the final Wiener filter output. We found that, even though the Wiener filter module is applied in feature space, some undesirable deconvolution artifacts, most noticeably ringings, are visible especially around image boundaries and object edges since the *FFT* operator within the Wiener deconvolution filter supposes circular periodicity of the input. To tackle this problem, an edgetaper operation [39] was implemented on the blurred input image to smooth out its boundaries which can considerably reduce the ringing artifacts in the final reconstructed sharp image.

We trained the network using the subset of FlyingThings3D [29] for 500 epochs with Adam optimizer [36] with exponential decay rates of the first momentum and second

momentum, respectively, set to $\beta_1 = 0.99$ and $\beta_2 = 0.999$, and a learning rate $L_r^{image} = 1e - 4$ which is halved after 250 epochs. The number of auto-encoders in the multi-scale feature refinement modules is set to 2 as in the original work of [28], the number of extracted feature maps from the blurry input is 16, and the batch size is set to 8. For the loss function, it was experimentally seen that *L1* norm leads to better reconstruction results than the ones obtained with *L2* norm.

$$\mathcal{L}_{image}(\theta, \theta^*) = \frac{1}{|T|} \sum_{\theta \in T} |\theta - \theta^*| \quad (11)$$

where θ and θ^* are, respectively, the reconstructed and the ground truth sharp images and $|T|$ is the number of images.

6 Experimental results

6.1 Experimental results on synthetic data

Monocular depth estimation on FlyingThings3D subset For this set of experiments the phase mask's trainable parameters are initialized to [$N = 1, \epsilon = 0.1$] and the number of Fresnel zones L is set to 7 as it was empirically observed that such value leads to a lower depth estimation error. In the case of noise-free RPSF-blurred inputs, the network achieved a RMSE of 0.392 on the test set. The corresponding learned phase mask parameters are [$N = 1, \epsilon = 0.92$]: the generated RPSF as well as the height map of the mask are shown in Fig. 5. The network learned a single-helix RPSF with high confinement parameter ϵ meaning that the peak is spread out across a large area. In all quantitative evaluations showed later on we use the same error and accuracy metrics used by prior works to assess the performance of monocular depth estimation solutions. They are the following:

- Error metrics:
 - Root Mean Squared Error (RMSE): it measures the difference between predicted and ground truth depth values.

$$RMSE(\theta, \theta^*) = \sqrt{\frac{1}{|T|} \sum_{\theta \in T} (\theta - \theta^*)^2} \quad (12)$$

- Relative error: It measures the average absolute deviation of predicted depth values with respect to ground truth values by which it is normalized.

$$Rel(\theta, \theta^*) = \frac{1}{|T|} \sum_{\theta \in T} \frac{|\theta - \theta^*|}{\theta^*} \quad (13)$$

- Log10 based error: Similar to mean absolute error, but computed on the logarithm of the depth values. This

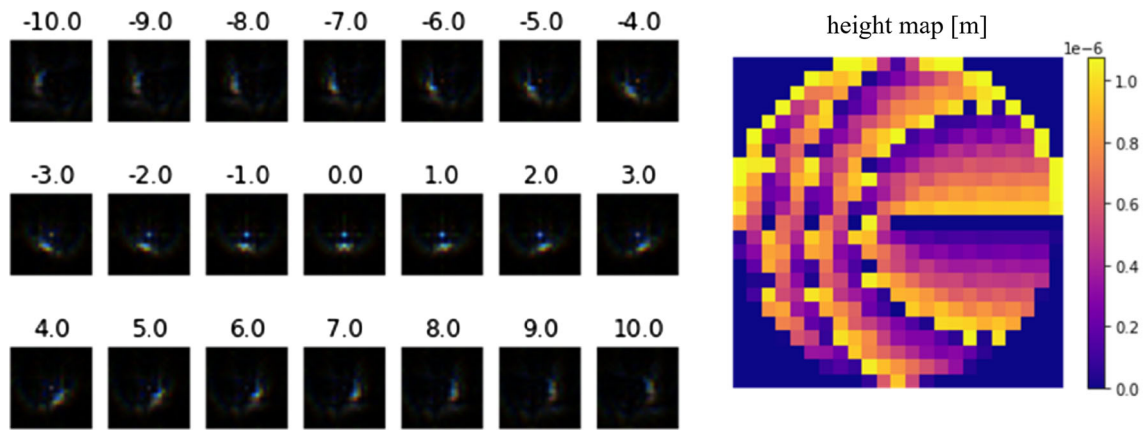


Fig. 5 Learned RPSF shape at different defocus planes (left), height map of the phase mask (right)

metric is often used to account for the scale-invariant nature of depth predictions.

$$\text{Log}_{10}(\bar{\theta}, \theta^*) = \frac{1}{|T|} \sum_{\theta \in T} |\log_{10}(\theta) - \log_{10}(\theta^*)| \quad (14)$$

- Accuracy metrics:

- Accuracy under a threshold: These metrics assess the percentage of pixels whose relative error falls below a certain threshold δ_i

$$\begin{aligned} \max\left(\frac{\theta}{\theta^*}, \frac{\theta^*}{\theta}\right) < \delta_1 \quad \delta_1 &= 1.25 \\ \max\left(\frac{\theta}{\theta^*}, \frac{\theta^*}{\theta}\right) < \delta_2 \quad \delta_2 &= 1.25^2 \\ \max\left(\frac{\theta}{\theta^*}, \frac{\theta^*}{\theta}\right) < \delta_3 \quad \delta_3 &= 1.25^3 \end{aligned} \quad (15)$$

where θ and θ^* are, respectively, the predicted and the ground truth depth values, and $|T|$ is the number of samples in the test set. All metrics are calculated between the ground truth and the predicted depth maps with pixel values in the range $]0, 10[$ meters.

Qualitative results are shown in Fig. 6 (additional scenes are shown in the supplementary material): notice how the network is able to predict accurate disparity maps with fine image details, e.g., the small leaves of the plant shown in the third row or the very fine parts of the headset shown in the last row. However, it becomes harder for the network to predict accurate object boundaries due to the blurring artifacts introduced by the RPSF especially when the blur kernel is larger than the image features. Some artifacts appear at object edges mainly due to the nature of the image formation model used to compose the RPSF-blurred images: the layered depth model used to render RPSF coded images does not accurately simulate the discontinuities around object edges as visible in the RPSF coded images (Fig. 6) due to the lack

of accurate occlusion modeling. Such issue can be addressed using more advanced and sophisticated blending and matting approaches, e.g., Pyramid-based blending [40], at the expense of much higher computation time and complexity for a minor gain in performance.

As pointed out in Sect. 4.2, we also introduced an accurate noise simulation model for a more realistic evaluation. To handle noise we further fine-tune the DEPTH-DNN on noisy images. In this case we achieve a RMSE of 0.712 on the test set compared to 0.392 achieved on noise-free data. This is of course expected but at the same time the robustness of the system in real world applications should be enhanced.

In the first two columns of Fig. 6, one can observe image quality degradation after simulating the noisy images where the color down-sampling by the CFA and quantization artifacts by the ADC unit are visible (zoomed-in). The predicted disparity maps from noisy inputs are shown in the last column. Even though a small performance degradation is noticeable on the noisy predictions, the fine-tuned network is still able to learn fairly accurate disparity maps.

Image restoration

The sharp all-in-focus images are recovered by IMAGE-DNN trained on the subset of FlyingThings3D [29]. Quantitative and qualitative results are shown in Table 1 and Fig. 7, respectively.

The simulated RPSF-coded images have a low mean PSNR of roughly 19 dB with respect to their sharp noise-free counterparts which were used as the ground truth images during training. As reported in Table 1, the mean PSNR of the recovered images increased by about 5.5 dB reaching 24.46 dB. Also, the Structural Similarity Index Measure (SSIM) achieved is 0.760 compared to 0.611 for the blurred and noisy images. Deblurring results from the traditional Wiener deconvolution filter [41] are also shown in Fig. 7: even if we apply the deconvolution process independently for each depth plane and generated the final result following Eq. 7, it results in a low quality image reconstruction with heavy ring-




Fig. 6 Qualitative results on RPSF-blurred images from the test set of FlyingThings3D [29] subset

Table 1 Quantitative results of the image deblurring model on the test set of FlyingThings3D [29] subset

	PSNR (dB)	SSIM
RPSF-coded images	19.01	0.611
Recovered images (Wiener)	18.75	0.449
Recovered images (Our approach)	24.46	0.760

Bold metrics indicate the best results

ing artifacts (see Fig. 7). This happens since the Wiener filter fails to handle the spatially variant blur producing significant ringing artifacts.

Figure 7 shows some recovered images along with the corresponding blurred inputs and the sharp all-in-focus ground truth. Upon visual inspection, one can notice that the model

successfully restores very fine image details and high frequency components, e.g., the small tree leaves shown in the first row, or the various background details present in the second row. Notice also how large regions with smooth as well as textured structures are recovered. The quantization noise and color down-sampling by the CFA make the task even more challenging resulting in some ringing artifacts on object edges.

Although an edgetaper [39] technique was used to limit such artifacts, few are still present in the recovered images, but are significantly reduced when comparing with the ones produced by the Wiener filter.




Fig. 7 Qualitative results of the image deblurring model on the test set of FlyingThings3D [29] subset

6.2 Experimental results on real data

NYUV2 depth dataset In order to compare our approach with state-of-the-art methods, the DEPTH-DNN is trained on a subset of NYUV2 indoor dataset [30] in an end-to-end fashion with the phase mask's height map. In this experiment, only 10 depth planes are considered in the layered depth presentation (Eq. 7) due to memory constraints. The simulated camera lens parameters are the following: $f/4.0$ with 4 mm aperture diameter and 16 mm focal length focusing at a distance of 5 m. The RGB images are directly convolved with a RPSF cube of the shape $10 \times 23 \times 23 \times 3$.

In the following evaluations, we applied the same crop used in competing works (e.g., [31]) and excluded the invalid pixels from the Kinect V1 sensor as done by competitors.

Figure 8 shows the learned RPSF and the corresponding phase mask. For this dataset, the learned RPSF is a double-helix [$N = 2, \epsilon = 0.99$] with two main side lobes that rotate with defocus. We argue that such behavior is mainly related to the characteristics of the training data: for more complex depth scenes, like in this case, the network tends to converge to higher number of peaks as it makes it easier to correlate between the rotation angle and the corresponding depth plane. Similar to the previous scenario, the network

also converges toward a high confinement parameter $\epsilon = 0.99$ producing more spread out peak regions.

Quantitative results on NYUV2 [31] test set are reported in Table 2. Coded-aperture based competing approaches [17, 18] optimized a free-form phase mask parameterized with a superposition of a set of Zernike polynomials [19, 20] using a U-Net [21] architecture. Besides using a more accurate camera model, differently from [17, 18], our approach learns only a few design parameters for the mask and simultaneously tackles the problem of image quality degradation. Table 2 shows quantitative results for the different error metrics used in [30]. Our approach outperforms the competing methods of [17, 18] in all but the last two accuracy metrics (where all top approaches including ours are very close to 1 making them not too significant). Note that for the more significant accuracy metric δ_1 (i.e., with the lowest threshold value), our approach achieves the highest score, even if we trained the network using a subset of 50k training samples which is less than half of the default split of 120k training samples [30] used by [17] and [18]. In particular, we achieve a significantly lower RMSE value of 0.267 which is down by, respectively, 0.166 and 0.115 from the ones achieved by [17] and [18] and is the lowest yet achieved on this challenging dataset for the task of monocular depth estimation.




Fig. 8 Learned RPSF shape at different defocus planes (left), the height map of the phase mask (right)

Table 2 Quantitative comparison with the state-of-the-art for monocular depth estimation task on NYUV2 test set [31]

Method	RMSE ↓	Rel ↓	Log10 ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
Eigen et al. [31]	0.641	0.158	–	0.769	0.950	0.988
Laina et al. [42]	0.573	0.127	0.055	0.811	0.953	0.988
Hao et al. [38]	0.555	0.127	0.053	0.841	0.966	0.991
DORN [43]	0.509	0.115	0.051	0.828	0.965	0.992
Qi et al. [44]	0.569	0.128	0.057	0.834	0.960	0.990
Yin et al. [45]	0.416	0.108	0.048	0.875	0.976	0.994
BTS [46]	0.392	0.110	0.047	0.885	0.978	0.994
DAV [47]	0.412	0.108	–	0.882	0.980	0.996
Alhashim et al. [48]	0.382	0.093	0.050	<u>0.932</u>	<u>0.989</u>	0.997
AdaBins [32]	0.364	0.103	<u>0.044</u>	0.903	0.984	0.997
DPT-Hybrid [49]	<u>0.357</u>	0.110	0.045	0.904	0.988	<u>0.998</u>
DeepOptics [17]	0.433	<u>0.087</u>	0.052	0.930	0.990	0.999
PhaseCam3D [18]	0.382	0.093	0.050	<u>0.932</u>	<u>0.989</u>	0.997
Ours	0.267	0.072	0.029	0.952	<u>0.989</u>	0.997

Bold metrics indicate the best results, underlined metrics indicate second best results

This performance gain is primarily related to the optimized PSF shape. The one obtained by [18] has a generic shape with no clear correlations between different defocus planes, while the one obtained by [17] has an elliptical shape with varying section which increases with depth. On the other hand, the RPSF shape produces a clear and simple correlation between the defocus plane and the corresponding angle of rotation of the main peaks, thus encoding robust depth cues within input images. This, thanks also to the small number of parameters to be learned, explain why the network achieves better depth estimation accuracy with significantly less training data.

Table 2 also presents quantitative results from the state-of-the-art that mainly used sharp all-in-focus images as input. Our approach outperforms all competing methods in all evaluation metrics except for δ_3 accuracy metric where all the top approaches including ours are extremely close. More specifically, we achieve substantially lower error metrics compared

to the competing approaches which used similar or larger training sets, e.g. Eigen et al. [31], Laina et al. [42], DORN [43], DAV [47], Alhashim et al. [48], AdaBins [32] and DPT-Hybrid [49], notice that the latter used a pre-trained model on a large combination of different datasets containing 1.4M samples and fine-tuned it on NYUV2 dataset [31]. The rest of the competing approaches used less training samples but at the cost of more complex models with pre-trained weights, e.g., Hao et al. [38] used a ResNet [50] backbone pre-trained on ImageNet [51]. Our model uses a smaller training set but the core network used for depth estimation is a simple U-Net [21] architecture with approximately 8.6M trainable parameters, the inference time needed to produce a single depth map with a spatial resolution of 256×256 pixels on an NVIDIA TITAN X is 15.85 msec and the GPU memory needed is 200 Mb. In contrast, AdaBins [32] for example, like most other competing methods used a more complex network architecture with approximately 78M trainable parameters, making




Fig. 9 Qualitative results from the proposed approach as well as those from AdaBins [32] on the test set of NYUV2 [31]

Table 3 Quantitative comparison with the state-of-the-art methods for monocular depth estimation task on SUNRGBD [33] test set

Method	RMSE ↓	Rel ↓	Log10 ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
Chen et al. [37]	0.494	0.166	0.071	0.757	0.943	<u>0.984</u>
Yin et al. [45]	0.541	0.183	0.082	0.696	0.912	0.973
BTS [46]	0.515	0.172	0.075	0.740	0.933	0.980
AdaBins [32]	<u>0.476</u>	<u>0.159</u>	<u>0.068</u>	<u>0.771</u>	<u>0.944</u>	0.983
Ours	0.335	0.114	0.034	0.937	0.981	0.992

Bold metrics indicate the best results, underlined metrics indicate second best results

it slower in both training and inference. In our case, the RPSF encodes strong and robust depth cues making it easier for a simple network to predict accurate depth maps compared to those using conventional RGB inputs.

A qualitative comparison with AdaBins [32] (trained with the same 50k samples as in this approach) is shown in Fig. 9 while further qualitative results are shown in the supplementary materials. Upon visual inspection, one can easily see that our approach produces more accurate and realistic depth maps with respect to the ground truth (the ground truth depth maps are inpainted for visualization purposes, in the supplementary material the ground truth data used for evaluation are shown). Due to the scarcity of reliable depth cues in single all-in-focus input images, Adabins [32] struggles to predict accurate and sometimes realistic depth values in a consistent manner and fails to predict correct values for images where

depth values span large ranges, e.g., the results shown in the two last columns in Fig. 9. Moreover, sometimes Adabins [32] produces erroneous depth predictions where the scene semantics are somehow confusing and the network fails to infer realistic values: such behavior exposes the main limitation of semantic-based approaches, as visible in the last column of Fig. 9 where the green carpet was misclassified. In contrast, our network consistently produces accurate depth maps for small and large depth ranges alike and is more agnostic to the scene's semantics. The main drawback of our method is that it sometimes fails to learn well defined object boundaries due to the blurring artifacts introduced by the RPSF kernel.

SUNRGBD dataset We evaluate the generalization capability of our model for the task of monocular depth estimation where DEPTH-DNN that was previously trained on the




Fig. 10 Qualitative results from our proposed method on the test set SUNRGBD [33]

Table 4 Quantitative results of the ablation experiments on the test set of FlyingThings3D [29]

Exp	Input	RMSE ↓	Rel ↓	Log10 ↓
1	All-in-focus	2.649	1.086	0.285
2	Fixed mask 1 [$N = 1, L = 5, \epsilon = 0.5$]	1.117	0.461	0.161
3	Fixed mask 2 [$N = 2, L = 5, \epsilon = 0.5$]	0.815	0.225	0.129
4	Learned mask [$L = 5$]	0.699	0.136	0.072

subset of NYUV2 [30] is evaluated using the test set of SUNRGBD [33] without any further fine-tuning. Quantitative and qualitative results are present in Table 3 and Fig. 10 (additional qualitative results are in the supplementary material). Metric values for competitors shown in Table 3 are taken from [32] where methods with publicly available pre-trained models on NYUV2 [30] have been evaluated on the SUNRGBD dataset.

As shown in Table 3, our approach outperforms the state-of-the-art in all evaluation metrics with a significant reduction in error metrics, particularly the RMSE (where it achieved 0.335 compared to 0.476 of the best competitor) and Log10 (0.034 against 0.068). Notice also that our approach achieves a δ_1 accuracy value (corresponding to the smallest threshold of 1.25) of 0.937, which is up by 0.166 compared to the best performing competitor that achieves 0.771. This indicates that our approach produces a higher percentage of accurately predicted depth values. Such results support the suitability of such engineered PSFs for depth acquisition applications enabling reliable and robust passive monocular depth estimation performance with real-time capabilities.

Figure 10 shows some prediction samples from SUNRGBD [33] test set, as in the previous case, the network is able to predict overall accurate depth maps but with higher mean RMSE compared to the test set of NYUV2 test set [31] which is expected due to the different statistical properties between the two datasets.

6.3 Ablation study

A number of experiments were carried out as an ablation study to assess the contribution of the various components in our framework. Quantitative results are shown in Table 4, while the qualitative ones are in the Supplementary Material. In all simulations of the ablation study, the network architecture as well as the training settings and hyper-parameters are the same as indicated in the previous section except for the usage of a simple Gaussian additive noise model instead of the full image formation procedure. Furthermore, we set the number of Fresnel zones in the mask design to $L = 5$.

The baseline is a U-Net trained on all-in-focus sharp images from the subset of FlyingThings3D [29]. The RMSE achieved in this first experiment is 2.649. In the second experiment, a coded aperture with a fixed phase mask design [$N = 1, L = 5, \epsilon = 0.5$] is used to blur the sharp input images with a depth dependent single-helix RPSF. In fact, this particular mask design is the one which was first introduced by [8]. The network trained with the fixed mask achieves a RMSE value of 1.117, i.e., down by 1.532 with respect to the baseline.

In the third experiment, the number of RPSF lobes is increased to $N = 2$. Both the number of Fresnel zones L and the confinement parameter ϵ are the same as in the second experiment. The network trained with such phase mask achieves even better RMSE value of 0.815 compared to the 1.117 achieved by the one in the previous experiment (see

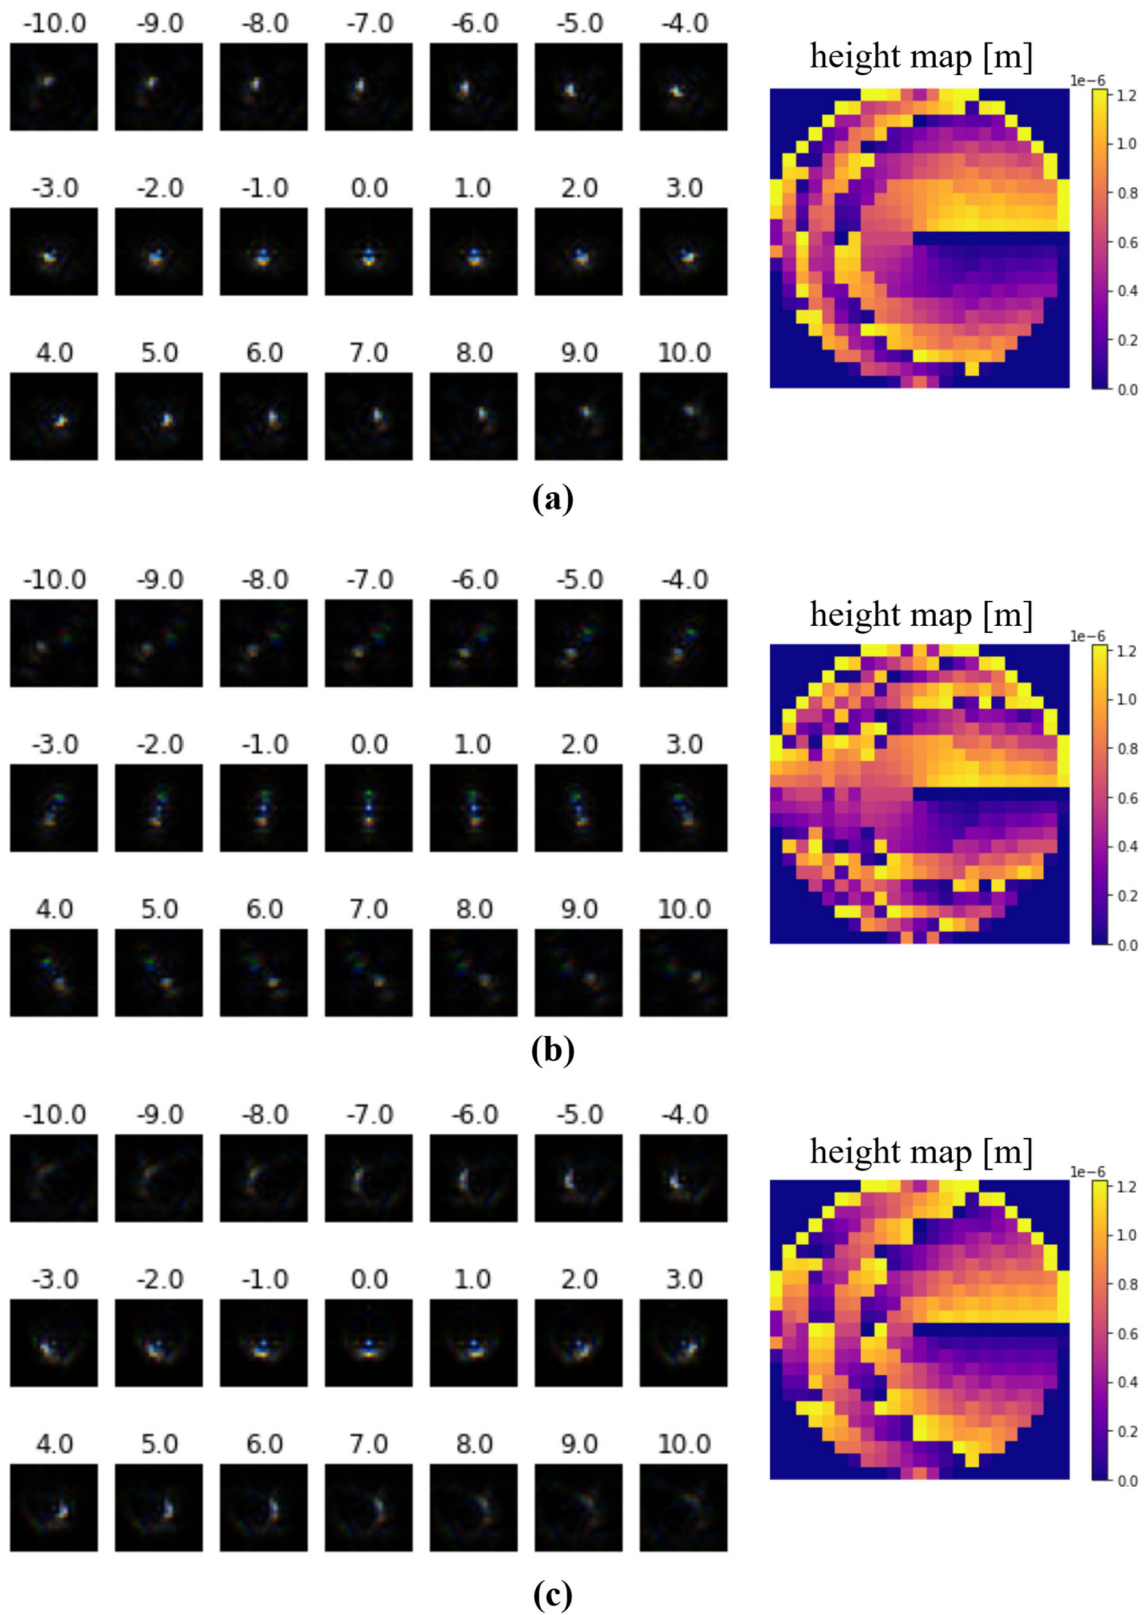


Fig. 11 The generated RPSFs for the ablation experiments: **a** Single-helix RPSF generated by the fixed mask 1 and its height map. **b** Double-helix RPSF generated by the fixed mask 2 and its height map. **c** Single-helix RPSF generated by the learned mask and its height map

Table 4). Such performance gain could be due to the more discriminative shape of the double-helix RPSF compared to the single-helix RPSF generated by the previous mask as the rotation can be easily noticed from a depth plane to the other. As shown in Fig. 11, the RPSF generated by the second mask has two main side peaks rotating counterclockwise as a function of defocus, the same rotation aspect can be observed in the RPSF shape generated by the first fixed mask (Fig. 11a) except that in this case only one main side lobe is present. It suggests that, for a fixed phase mask, a double-helix RPSF conveys more discriminative depth cues than a single-helix one.

In the fourth and last experiment, the phase mask's trainable parameters N and ϵ are jointly optimized with the weights of the network, the number of Fresnel-zones is fixed to $L = 5$. As expected, the network was able to outperform the baseline as well as the ones trained with fixed masks, reaching a RMSE of 0.699. Notice how the other two metrics shown in the table confirm the evaluation done on RMSE. Furthermore, results can be improved by using more complex backbone networks, as an example using a RESUNet [52] allows to improve the RMSE to 0.572 but at the price of a four times increase in the computation time. Since real-time applications are a key target for the approach we decided to stick on the baseline model.

The learned phase mask parameters are [$N = 1, \epsilon = 0.91$] meaning that the RPSF (shown in Fig. 11c) has a single side lobe that rotate with defocus. Notice that also the confinement parameter ϵ is high resulting in a more spread out lobe compared to the one generated by the fixed mask. It is therefore clear that a joint optimization approach helps the network to effectively learn the correlations between the rotation angle of the PSF and the corresponding depth plane leading to better estimation accuracy.

7 Conclusion

In this paper, we presented a novel computational camera model where an end-to-end learning framework is proposed for the joint optimization of camera optics and image processing algorithms for the tasks of monocular depth estimation from diffracted rotation. The learned phase mask generates multi-order helix rotating PSFs as a function of defocus, encoding strong depth cues within single 2D images and enabling reliable and accurate depth estimation. Experimental results confirmed the capability of the proposed model to outperform existing methods in the task of monocular depth estimation and to generalize well beyond the training environment. The depth estimation model complexity is significantly reduced compared to the state-of-the-art due to the 3D cues encoded by the RPSF, making it suitable for real-time applications without compromising accuracy. Finally,

the sharp all-in-focus images are also recovered through a dedicated non-blind and nonuniform image deblurring module.

Further research will focus on the fabrication of the phase mask via photo-lithography which will be mounted on the back side of the camera's aperture, thus adding depth estimation capabilities to a standard RGB cameras. This could enable multiple downstream tasks such as 3D object detection, tracking and SLAM, even if the impact of the image blurring due to the modified optics on these tasks need to be considered and appropriately tackled. Post capture image manipulation will also be considered.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s00371-023-03147-8>.

Funding Open access funding provided by Università degli Studi di Padova within the CRUI-CARE Agreement. The work was supported by Sony Europe B.V.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Greengard, A., Schechner, Y.Y., Piestun, R.: Depth from diffracted rotation. *Opt. Lett.* **31**(2), 181–183 (2006)
2. Allen, L., Beijersbergen, M.W., Spreeuw, R., Woerdman, J.: Orbital angular momentum of light and the transformation of Laguerre–Gaussian laser modes. *Phys. Rev. A* **45**(11), 8185 (1992)
3. Piestun, R., Schechner, Y.Y., Shamir, J.: Propagation-invariant wave fields with finite energy. *JOSA A* **17**(2), 294–303 (2000)
4. Pavani, S.R.P., Piestun, R.: High-efficiency rotating point spread functions. *Opt. Express* **16**(5), 3484–3489 (2008)
5. Shechtman, Y., Sahl, S.J., Backer, A.S., Moerner, W.: Optimal point spread function design for 3d imaging. *Phys. Rev. Lett.* **113**(13), 133902 (2014)
6. Kotlyar, V.V., Almazov, A.A., Khonina, S.N., Soifer, V.A., Elfstrom, H., Turunen, J.: Generation of phase singularity through diffracting a plane or gaussian beam by a spiral phase plate. *JOSA A* **22**(5), 849–861 (2005)
7. Kotlyar, V.V., Kovalev, A.A., Khonina, S.N., Skidanov, R.V., Soifer, V.A., Elfstrom, H., Tossavainen, N., Turunen, J.: Diffraction of

- conic and gaussian beams by a spiral phase plate. *Appl. Opt.* **45**(12), 2656–2665 (2006)
8. Prasad, S.: Rotating point spread function via pupil-phase engineering. *Opt. Lett.* **38**(4), 585–587 (2013)
 9. Berlich, R., Stallinga, S.: High-order-helix point spread functions for monocular three-dimensional imaging with superior aberration robustness. *Opt. Express* **26**(4), 4873–4891 (2018)
 10. Kumar, R.: Three-dimensional imaging using a novel rotating point spread function imager, The University of New Mexico (2015)
 11. Levin, A., Fergus, R., Durand, F., Freeman, W.T.: Image and depth from a conventional camera with a coded aperture. *ACM Trans Graph (TOG)* **26**(3), 70–526 (2007)
 12. Zhou, C., Lin, S., Nayar, S.K.: Coded aperture pairs for depth from defocus and defocus deblurring. *Int. J. Comput. Vision* **93**(1), 53–72 (2011)
 13. Kumar, R., Prasad, S.: Psf rotation with changing defocus and applications to 3d imaging for space situational awareness, In: Proceedings of the 2013 AMOS Technical Conference, Maui, HI (2013)
 14. Roider, C., Jesacher, A., Bernet, S., Ritsch-Marte, M.: Axial super-localisation using rotating point spread functions shaped by polarisation-dependent phase modulation. *Opt. Express* **22**(4), 4029–4037 (2014)
 15. Berlich, R., Bräuer, A., Stallinga, S.: Single shot three-dimensional imaging using an engineered point spread function. *Opt. Express* **24**(6), 5946–5960 (2016)
 16. Haim, H., Elmalem, S., Giryas, R., Bronstein, A.M., Marom, E.: Depth estimation from a single image using deep learned phase coded mask. *IEEE Trans. Comput. Imag.* **4**(3), 298–310 (2018)
 17. Chang, J., Wetzstein, G.: Deep optics for monocular depth estimation and 3d object detection, In: Proceedings of the IEEE/CVF international conference on computer vision, pp. 10193–10202 (2019)
 18. Wu, Y., Boominathan, V., Chen, H., Sankaranarayanan, A., Veeraraghavan, A.: Phasecam3d-learning phase masks for passive single view depth estimation, In: 2019 IEEE International Conference on Computational Photography (ICCP), IEEE, pp. 1–12 (2019)
 19. Iskander, D.R., Collins, M.J., Davis, B.: Optimal modeling of corneal surfaces with zernike polynomials. *IEEE Trans. Biomed. Eng.* **48**(1), 87–95 (2001)
 20. Born, M., Wolf, E.: Principles of optics: electromagnetic theory of propagation, interference and diffraction of light. Elsevier, Amsterdam (2013)
 21. Ronneberger, O., Fischer, P., Brox, T.: U-net Convolutional networks for biomedical image segmentation. In: International conference on medical image computing and computer-assisted intervention, pp. 234–241. Springer, Berlin (2015)
 22. Anger, J., Facciolo, G., Delbracio, M.: Modeling realistic degradations in non-blind deconvolution, In: 2018 25th IEEE International conference on image processing (ICIP), IEEE, pp. 978–982 (2018)
 23. Cho, S., Lee, S.: Fast motion deblurring, In: ACM SIGGRAPH Asia 2009 papers, pp. 1–8 (2009)
 24. Gong, D., Zhang, Z., Shi, Q., van den Hengel, A., Shen, C., Zhang, Y.: Learning deep gradient descent optimization for image deconvolution. *IEEE Trans. Neural Netw. Learn. Syst.* **31**(12), 5468–5482 (2020)
 25. Krishnan, D., Fergus, R.: Fast image deconvolution using hyper-Laplacian priors. *Adv. Neural. Inf. Process. Syst.* **22**, 1033–1041 (2009)
 26. Xu, L., Ren, J.S., Liu, C., Jia, J.: Deep convolutional neural network for image deconvolution. *Adv. Neural. Inf. Process. Syst.* **27**, 1790–1798 (2014)
 27. Schuler, C.J., Christopher Burger, H., Harmeling, S., Scholkopf, B.: A machine learning approach for non-blind image deconvolution, In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1067–1074 (2013)
 28. Dong, J., Roth, S., Schiele, B.: Deep wiener deconvolution: wiener meets deep learning for image deblurring, arXiv preprint [arXiv:2103.09962](https://arxiv.org/abs/2103.09962)
 29. Mayer, N., Ilg, E., Häusser, P., Fischer, P., Cremers, D., Dosovitskiy, A., Brox, T.: A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation, In: IEEE International conference on computer vision and pattern recognition (CVPR), 2016, [arXiv:1512.02134](https://arxiv.org/abs/1512.02134). <http://lmb.informatik.uni-freiburg.de/Publications/2016/MIFDB16>
 30. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network, arXiv preprint [arXiv:1406.2283](https://arxiv.org/abs/1406.2283)
 31. Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture, In: Proceedings of the IEEE international conference on computer vision, pp. 2650–2658 (2015)
 32. Bhat, S.F., Alhashim, I., Wonka, P.: Adabins: depth estimation using adaptive bins, arXiv preprint [arXiv:2011.14141](https://arxiv.org/abs/2011.14141)
 33. Song, S., Lichtenberg, S.P., Xiao, J.: Sun rgb-d: A rgb-d scene understanding benchmark suite, In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 567–576 (2015)
 34. Malvar, H.S., He, L.-w., Cutler, R.: High-quality linear interpolation for demosaicing of bayer-patterned color images, In: 2004 IEEE International conference on acoustics, speech, and signal processing, Vol. 3, IEEE, pp. iii–485 (2004)
 35. M. A. et al., TensorFlow: large-scale machine learning on heterogeneous systems, software available from tensorflow.org (2015). <http://tensorflow.org/>
 36. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization, arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
 37. Chen, X., Chen, X., Zha, Z.-J.: Structure-aware residual pyramid network for monocular depth estimation, arXiv preprint [arXiv:1907.06023](https://arxiv.org/abs/1907.06023)
 38. Hao, Z., Li, Y., You, S., Lu, F.: Detail preserving depth estimation from a single image using attention guided networks, In: 2018 International Conference on 3D Vision (3DV), IEEE, pp. 304–313 (2018)
 39. Reeves, S.J.: Fast image restoration without boundary artifacts. *IEEE Trans. Image Process.* **14**(10), 1448–1453 (2005)
 40. Kraus, M., Strengert, M.: Depth-of-field rendering by pyramidal image processing, In: Computer graphics forum, Vol. 26, Wiley Online Library, pp. 645–654 (2007)
 41. Wiener, N., et al.: Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications, vol. 8. MIT press, Cambridge (1964)
 42. Laina, I., Rupprecht, C., Belagiannis, V., Tombari, F., Navab, N.: Deeper depth prediction with fully convolutional residual networks, In: 2016 Fourth international conference on 3D vision (3DV). IEEE, 239–248 (2016)
 43. Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation, In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2002–2011 (2018)
 44. Qi, X., Liao, R., Liu, Z., Urtasun, R., Jia, J.: Geonet: Geometric neural network for joint depth and surface normal estimation, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 283–291 (2018)
 45. Yin, W., Liu, Y., Shen, C., Yan, Y.: Enforcing geometric constraints of virtual normal for depth prediction, In: Proceedings of the IEEE/CVF international conference on computer vision, pp. 5684–5693 (2019)

46. Lee, J.H., Han, M.-K., Ko, D.W., Suh, I.H.: From big to small: multi-scale local planar guidance for monocular depth estimation, arXiv preprint [arXiv:1907.10326](https://arxiv.org/abs/1907.10326)
47. Huynh, L., Nguyen-Ha, P., Matas, J., Rahtu, E., Heikkilä, J.: Guiding monocular depth estimation using depth-attention volume, In: European conference on computer vision, Springer, pp. 581–597 (2020)
48. Alhashim, I., Wonka, P.: High quality monocular depth estimation via transfer learning, arXiv preprint [arXiv:1812.11941](https://arxiv.org/abs/1812.11941)
49. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction, arXiv preprint [arXiv:2103.13413](https://arxiv.org/abs/2103.13413)
50. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition, In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778 (2016)
51. Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database, In: 2009 IEEE Conference on computer vision and pattern recognition. IEEE, 248–255 (2009)
52. Diakogiannis, F.I., Waldner, F., Caccetta, P., Wu, C.: Resunet-a: a deep learning framework for semantic segmentation of remotely sensed data. *ISPRS J. Photogramm. Remote. Sens.* **162**, 94–114 (2020)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Mazen Mel received an Engineering degree in Telecommunications from the Higher School of Communications of Tunis in 2019 and a M.Sc. degree in Telecommunication from the University of Padova in 2021. He is currently a Ph.D. student at the Department of Information Engineering of the University of Padova. His research focuses on computational photography, in particular hyperspectral imaging, depth estimation, and depth of field extension.



Muhammad Siddiqui received a M.Sc. degree in Information Technology at the University of Stuttgart in 2004. Currently he is working as Senior Engineer at Sony Europe B.V. in Stuttgart. His research interests include computational imaging, image/video signal processing as well as compressive sensing.



Pietro Zanuttigh received the M.Sc. degree in Computer Engineering from the University of Padova in 2003 where he also got the Ph.D. degree in 2007. Currently he is an associate professor at the Department of Information Engineering. His research activity focuses on image and 3D data processing and analysis, in particular, ToF sensors data processing, domain adaptation and incremental learning in semantic segmentation, and hand gesture recognition.