Quantization-aware Deep Optics for Diffractive Snapshot Hyperspectral Imaging

Lingen Li Lizhi Wang* Weitao Song Lei Zhang
Beijing Institute of Technology
{lingenli, wanglizhi, swt, leizhang}@bit.edu.cn

Zhiwei Xiong
University of Science and Technology of China
zwxiong@ustc.edu.cn

Hua Huang
Beijing Normal University
huahuang@bnu.edu.cn

Abstract

Diffractive snapshot hyperspectral imaging based on the deep optics framework has been striving to capture the spectral images of dynamic scenes. However, existing deep optics frameworks all suffer from the mismatch between the optical hardware and the reconstruction algorithm due to the quantization operation in the diffractive optical element (DOE) fabrication, leading to the limited performance of hyperspectral imaging in practice. In this paper, we propose the quantization-aware deep optics for diffractive snapshot hyperspectral imaging. Our key observation is that common lithography techniques used in fabricating DOEs need to quantize the DOE height map to a few levels, and can freely set the height for each level. Therefore, we propose to integrate the quantization operation into the DOE height map optimization and design an adaptive mechanism to adjust the physical height of each quantization level. According to the optimization, we fabricate the quantized DOE directly and build a diffractive hyperspectral snapshot imaging system. Our method develops the deep optics framework to be more practical through the awareness of and adaptation to the quantization operation of the DOE physical structure, making the fabricated DOE and the reconstruction algorithm match each other systematically. Extensive synthetic simulation and real hardware experiments validate the superior performance of our method.

1. Introduction

Hyperspectral imaging has been beneficial to a plethora of sensing applications, from the fundamental research fields, *e.g.*, biomedical inspection, material analysis, and environmental monitoring [1, 4, 9, 27], to computer vision applications, *e.g.* appearance acquisition, face recognition, and object tracking [13, 32, 46]. Conventional hyperspectral imaging systems scan scenes along either the spatial or the spectral dimensions and thus require multiple sensor exposures to capture a full hyperspectral image [21]. These systems are unsuitable for measuring dynamic scenes, making snapshot hyperspectral imaging with just one sensor exposure one of the most attractive solutions.

Various snapshot hyperspectral imaging systems have been developed for capturing hyperspectral images of dynamic scenes [3, 11]. Classic methods based on geometrical optics mainly consist of a series of refractive or reflective optical elements [10, 25, 28, 41], being large in form factor and thus suffering from the system complexity and calibration difficulty. To overcome the limitations, diffractive snapshot hyperspectral imaging systems have been developed by replacing the geometrical optical elements with a thin diffractive optical element (DOE) [6, 18, 22, 33]. Fundamentally, diffractive snapshot hyperspectral imaging needs to design the DOE height map for hyperspectral image encoding and the reconstruction algorithm for hyperspectral image decoding. Previous methods try to design the DOE height map with hand-crafted or heuristic knowledge on the point spread function (PSF) but isolate the DOE hardware from the reconstruction algorithm [22]. Recently, the deep optics framework has been introduced with the core idea of an end-to-end optimization of optical hardware and the reconstruction algorithm [12, 20, 29, 36, 37], which has been the most promising method in diffractive snapshot hyperspectral imaging [5, 15].

The bottleneck in the deep optics framework is that the DOE height map optimization does not model the physical quantization in DOE fabrication. Existing deep optics frameworks all employ full precision (generally 32-bit float) in DOE height map optimization and assume the same full precision in DOE fabrication. However, common

*Corresponding author: Lizhi Wang.

lithography techniques used in fabricating physical DOEs need to configure the quantization levels of the DOE height map [30]. Mature lithography techniques usually support no more than 16 quantization levels and only 4 quantization levels if high lithography stability is required. In this case, the DOE height map optimized with full precision has to be quantized during the physical fabrication. Thus, the fabricated DOE would deviate from the optimized DOE due to the quantization operation, which breaks the link and leads to a mismatch between the optical hardware and the reconstruction algorithm, reducing the energy of deep optics and limiting the performance of the diffractive snapshot hyperspectral imaging.

In this paper, we propose the quantization-aware deep optics (QDO) for diffractive snapshot hyperspectral imaging. Specifically, We explicitly model the quantization operation during the DOE height map optimization, which is conducted together with the optimization of reconstruction algorithm, to bridge the gap between the DOE optimization and the DOE fabrication. By leveraging the fact that the DOE height map of each quantization level can be freely set to be an arbitrary value within a specific range, we further propose an adaptive quantization-aware mechanism by adjusting the physical height of each level during DOE height map optimization. Finally, we fabricate the DOE with the height map quantized according to the optimization and build a diffractive hyperspectral snapshot imaging system.

Our method develops the deep optics framework to be more practical through the awareness of and adaptation to the quantization operation of the DOE physical structure, making the fabricated DOE and the reconstruction algorithm match each other systematically. Extensive synthetic simulation and real hardware experiments validate the superior performance of our method.

Our contributions are summarized as follows:

- We propose a quantization-aware deep optics model by jointly optimizing the quantized DOE and the reconstruction algorithm for diffractive snapshot hyperspectral imaging.
- We propose an adaptive mechanism by adjusting the physical height of each quantization level during DOE height map optimization to approximate the oracle performance with impossible full-precision fabrication.
- We fabricate the quantized DOE optimized by the proposed method and build a diffractive snapshot hyperspectral imaging system.

2. Related Work

Snapshot Hyperspectral Imaging. Snapshot hyperspectral imaging needs only one 2D exposure to capture the encoded information of the 3D hyperspectral images, which has a distinct advantage over conventional hyperspectral imaging systems in capturing dynamic scenes. Based on compressive sensing theory [14], coded aperture snapshot spectral imagers (CASSI) is one of the representative solutions that can efficiently encode the hyperspectral image into one sensor image [3, 19, 42–45, 48]. However, CASSI uses a series of refractive or reflective optical elements as the optical encoder and suffers from system complexity and calibration difficulty. Thus, actual imaging applications are limited to laboratory environments. With the development of computational photography, diffractive hyperspectral imaging systems have been introduced that employ diffractive optic elements to reduce the physical size and complexity. Jeon *et al*. engineer the DOE with the heuristic idea to generate an anisotropic shape of the spectrally varying PSF [22]. Dun *et al*. learn a rotationally symmetric DOE with concentric ring decomposition [15]. Baek *et al*. extend the diffractive snapshot hyperspectral imaging by integrating with the depth information [5]. Diffractive snapshot hyperspectral imaging systems circumvent the need for many optical elements, allowing casual users to capture hyperspectral images. Our work inherits the advantages of diffractive hyperspectral imaging and tries to fix the deviation between theoretic modeling and physical fabrication.

Deep Optics. The deep optics framework, which jointly optimizes an optical encoder and computational decoder in an end-to-end manner, has been widely explored in recent years [5, 12, 15, 20, 29, 36–40]. Sitzmann *et al*. apply the deep optics framework on extended depth of field and super-resolution imaging [36]. Chang and Wetzstein use deep optics for monocular depth estimation and 3D object detection [12]. Ikoma *et al*. propose a more accurate occlusion-aware image formation model for deep optics in monocular depth estimation [20]. Besides, the deep optics has been applied in hyperspectral imaging [5, 15].

However, current deep optics models do not consider the impact of fabrication quantization in the simulation model, regarding the height map as a smooth structure (32-bit float). This problem causes the mismatch between the optimized and the fabricated DOE, which degrades reconstruction quality in physical systems. In contrast, our quantization-aware models fix this mismatch by directly modeling a quantized DOE.

Quantization-aware Training. Quantization-aware training is demanded in low-precision network training. Unlike the post-training quantization method, the forward and backward propagation are performed on models with quantized weights [16, 23]. The critical challenge of quantization-aware training is to approximate the gradient of non-differentiable operations. The common approach Straight Through Estimator (STE) [8] approximates every non-differentiable operation as an identity function. Although STE performs well in practice [34, 50], its training

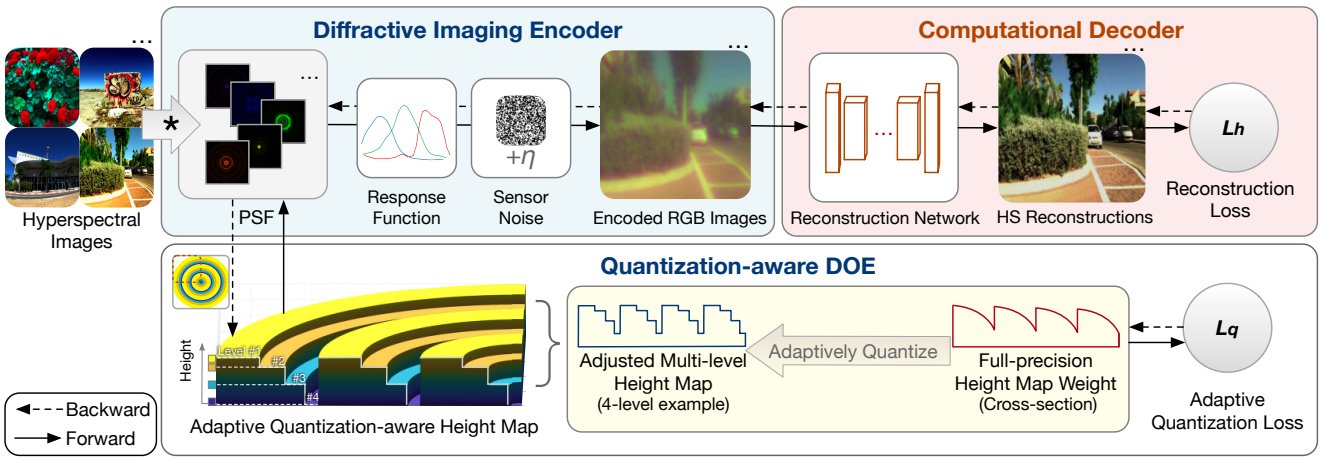

Figure 1. The overview of the proposed quantization-aware deep optics for diffractive hyperspectral snapshot imaging. The quantization-aware DOE produces a height map in multi-level style, and the diffractive imaging model uses the quantized height map to encode the hyperspectral image to one RGB image. The computational decoder takes the encoded RGB image as input and reconstructs the original hyperspectral image. All the pipeline, including the quantization-aware DOE is differentiable so that we can train the entire model in an end-to-end manner.

process is hard to be interpreted mathematically and unstable in certain situations [47]. Other methods solve this problem by avoiding direct use of non-differentiable operations [7, 24, 26, 31]. Particularly, Liu *et al*. [26] propose the alpha-blending approach using a combination of a full-precision branch and a quantized branch to represent the quantized weight, which allows the gradient update to the full-precision weight in backward propagation and finally transit the weight to a quantized result.

Our method integrates the quantization-aware training method alpha-blending with the deep optics framework to train a quantized DOE height map. Unlike the purpose of low-precision training that compresses models, what we need is a height profile that can be directly used for DOE fabrication. Thus, we further propose an adaptive mechanism to the quantization-aware training approach to reduce the quantization error, which promotes the quantized DOE to be more suitable for this joint optimization task.

3. Diffractive Imaging Model

Diffractive snapshot hyperspectral imaging is striving to capture spectral images of dynamic scenes. We propose a new imaging framework, QDO, for diffractive snapshot hyperspectral imaging, as shown in Fig. 1. This section describes the diffractive imaging model as the foundation of QDO.

3.1. Point Spread Function

The imaging system consists of a DOE for optical encoding and a bare RGB sensor for capturing the encoded image. We build the PSF of the system based on Fourier optics [17].

Fig. 2 illustrates the process of the wave field emitted from a point source light propagating through the DOE to the sensor plane. Suppose the point source with wavelength λ is located at a distance d from the DOE ($d \gg \lambda$), the wave field at the position (x, y) of the DOE can be formulated as:

$$U_0(x, y, \lambda) = e^{i\frac{2\pi}{\lambda}\frac{x^2+y^2}{d}}, \tag{1}$$

where i is the imaginary unit.

Once the wave field passes through the DOE, A phase delay $\phi(x, y)$ is introduced. Then, the wave field $U_1(x, y, \lambda)$ can be formulated as

$$U_1(x, y, \lambda) = A(x, y)U_0(x, y, \lambda)e^{i\frac{2\pi}{\lambda}\phi(x,y,\lambda)}, \tag{2}$$

where $A(x, y)$ is the optical aperture of the system. The phase delay $\phi(x, y)$ is determined by the height map of the DOE $H(x, y)$ as:

$$\phi(x, y, \lambda) = (n_\lambda - 1)H(x, y). \tag{3}$$

where n_λ is the refractive index of the DOE material at wavelength λ. Generally, the height map of the DOE $H(x, y)$ has a rotationally symmetric parameterization to reduce the computational complexity [15].

When the wave field reaches the sensor plane at depth z, the wave field $U_2(x, y, \lambda)$ can be obtained from $U_1(x, y, \lambda)$ by Fresnel diffraction law.

$$U_2(x, y, \lambda) = \mathcal{F}^{-1}\left\{\mathcal{F}\{U_1\}\, e^{i\frac{2\pi}{\lambda}z} e^{-i\pi\lambda z\left(f_x^2+f_y^2\right)}\right\}, \tag{4}$$

where f_x and f_y are the frequency variables of x and y respectively, and $\mathcal{F}$ denotes the Fourier transform.

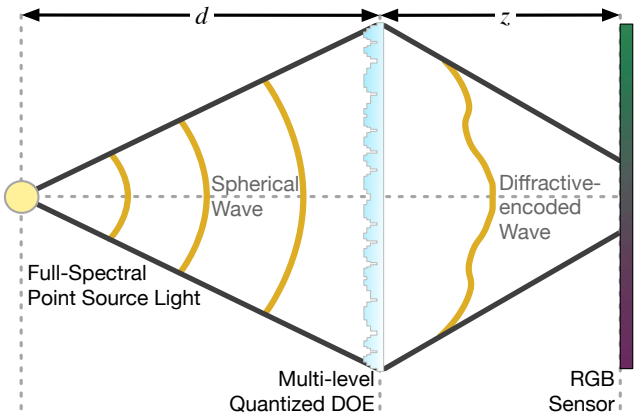

Figure 2. The propagation of the wave field from the source through the quantized DOE to the sensor plane.

The PSF P is the intensity of the squared value of the wave field $U_2(x, y, \lambda)$ as:

$$P(x, y, \lambda) \propto |U_2(x, y, \lambda)|^2 \tag{5}$$

3.2. Sensor Model

After formulating the PSF, we can use the system to capture images. The original hyperspectral images $I(x, y, \lambda)$ will be firstly modulated by the PSF as a convolution:

$$I'(x, y, \lambda) = P(x, y, \lambda) \otimes I(x, y, \lambda), \tag{6}$$

where $\otimes$ is the convolution operator.

Then the hyperspectral image is captured by the sensor with a spectral response function $R_c(\lambda)$ for each wavelength and becomes an RGB image $I_{c\in\{R,G,B\}}$. This imaging process usually introduces some sensor noise.

$$I_{c\in\{R,G,B\}}(x, y) = \int_{\lambda_0}^{\lambda_1} I'(x, y, \lambda) R_c(\lambda) d\lambda + \eta, \tag{7}$$

where λ_0 is the minimum wavelength, λ_1 is the maximum wavelength, and η is the sensor noise.

4. Quantization-aware DOE

From the diffractive imaging model in Sec. 3, we can see that the height map of the DOE $H(x, y)$ determines the PSF of the systems. Thus, great efforts have been made to optimized the DOE height map to realize customized encoding [15, 37]. Here, we want to emphasize that existing deep optics frameworks assume full-precision height map in the PSF formulation, which is ideal but impossible in practice. In this section, we introduce the practical QDO, which models the quantization operation in the PSF formulation to obey the physics in the DOE fabrication. Based on the QDO model, we further propose an adaptive mechanism to decide the height map for each quantization level. We name the QDO mode with the proposed adaptive mechanism as "QDO+A".

4.1. Quantization-aware Model

Here, we introduce the quantization-aware model that considers the quantization operation during the DOE optimization. We leverage the alpha-blending method [26] to train the quantized DOE. Assume we can obtain a full-precision height map H_f through direct optimization like previous methods [5, 15]. Then the quantization-aware height map H_q is a weighted summation of the quantized branch and the full-precision branch:

$$H_q = \alpha \times Q(H_f) + (1 - \alpha) \times H_f, \tag{8}$$

where $Q(\cdot)$ is the quantization function and α is the blending parameter.

The quantization function $Q(\cdot)$ is used to quantize the full-precision height map into L evenly divided levels $h_{l\in\{1,2,\cdots,L\}}$ within the max physical height h_{max}:

$$Q(H) = h_{max} \times \frac{\lfloor (L-1) \times H_f / h_{max}) + 0.5 \rfloor}{L - 1} \tag{9}$$

The blending parameter α increases with the training step s as:

$$\alpha(s) = \begin{cases} 0 & s \le T_0 \\ 1 - \left(\frac{T_1 - s}{T_1 - T_0}\right)^3 & T_0 < s \le T_1 \\ 1 & T_1 < s \end{cases}, \tag{10}$$

where T_0 and T_1 are the indicators of the training step when the blending starts and ends respectively. At training steps before T_0, α is 0 and the height map H_q is full-precision. During training steps between T_0 and T_1, the cubic function changes α from 0 to 1, and the quantization-aware height map H_q is transiting from the full-precision version to the quantized version. After step T_1, the α becomes 1, and the height map H_q is completely quantized.

4.2. Adaptive Mechanism

The full precision height map is ideal in the deep optics framework, so the quantized DOE under the constraint of quantization level should as much as possible approximate the oracle performance of the full precision optimization. Based on the observation that DOE height map of each quantization level can be freely set to be arbitrary value within a specific range, we further propose an adaptive quantization-aware mechanism by adjusting the physical height of each level during DOE height map optimization, which can effectively decrease the quantization deviation from the full precision height map.

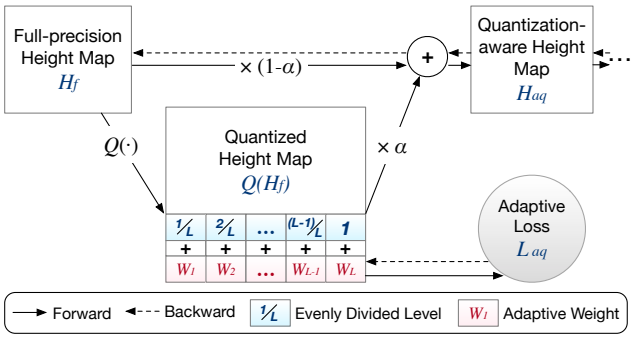

Figure 3. The adaptive mechanism for the quantization-aware training model. The parameter α controls the blending process of the height map from a full-precision structure to a quantized level structure (4-level as the example here). The trainable adaptive weights will adjust the physical height of each quantization level during training.

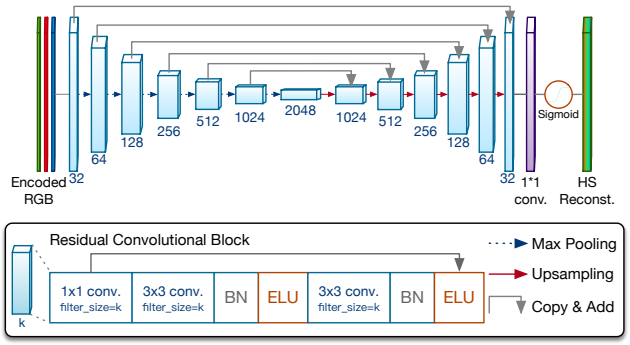

Figure 4. The structure of reconstruction network Res-UNet in our implementation. There are 13 residual convolution blocks in total, of which 6 blocks perform max pooling on their output, and the other 6 blocks upsample their input. Long connections are added across the blocks with the same filter size. Its output will be clipped to $(0, 1)$ by the final Sigmoid activation.

With the adaptive mechanism, the quantization-aware height map H_{aq} can be formulated as:

$$H_{aq} = \alpha \times F(Q(H_f)) + (1 - \alpha) \times H_f, \tag{11}$$

where $F(\cdot)$ is the adaptive operator performed on the evenly quantized height map $Q(H_f)$. Fig. 3 shows the detailes of the adaptive mechanism. Specifically, The adaptive operator adds a trainable weight $W_{l\in\{1,2,...,L\}}$ on the quantized height map of the DOE in each quantization level as:

$$F(Q(H_f)) = Q(H_f) + W_l \tag{12}$$

where l indicates the specific quantization level.

The physical height of each level is adjusted by training $W_{l\in\{1,2,...,L\}}$ to minimize the mean square error between the quantized height map and the full-precision height map as:

$$\mathcal{L}_{aq} = \frac{1}{J}\|F(Q(H_f)) - H_f\|_2^2, \tag{13}$$

where J indicates the pixel count of the height map. This objective is treated as a part of the training loss.

5. Hyperspectral Image Reconstruction

Once the hyperspectral image is encoded into an RGB image by following the diffractive imaging model, it needs an computational decoder to reconstruct hyperspectral image from the encoded image $I_{c\in\{R,G,B\}}$. We adopt the Res-UNet [49], a variant of the U-Net [35], as the decoder. As shown in Fig. 4, we configure six layers in both downsampling and upsampling stages and a middle layer between the two stages. Each layer is a residual convolutional block and ELU is used as the activation for each layer. We also add an extra convolution layer with Sigmoid activation to limit the output values between 0 and 1.

The training loss function $\mathcal{L}$ consists of reconstruction loss $\mathcal{L}_h$ on the hyperspectral image reconstruction, adaptive quantization loss $\mathcal{L}_{aq}$ on the DOE height map, and L_2 regularization on network weights:

$$\mathcal{L} = \mathcal{L}_h + \beta\mathcal{L}_{aq} + \gamma\|\omega\|_2^2, \tag{14}$$

where β and γ are scaling parameters, which are set to 10^{-2} and 10^{-4} respectively, and ω denotes weights of the reconstruction network. The reconstruction loss $\mathcal{L}_h$ evaluates the mean absolute error between the reconstructed hyperspectral image $\hat{I}$ and the ground truth I:

$$\mathcal{L}_h = \frac{1}{K}\|\hat{I} - I\|_1, \tag{15}$$

where K denotes the pixel count of the image.

The DOE height map and the reconstruction network are trained in an end-to-end manner. We train the whole model for 50 epochs with a batch size of 4 using Adam optimizer. The initial learning rate is 0.01 for the DOE part and 0.001 for the network part, with decays rates of 80% for every epoch. The quantization-aware blending starts at epoch 5 and ends at epoch 40. We use the NVIDIA GeForce RTX 3090 GPU as the computation platform, and all models are implemented using TensorFlow 2. We have released the code at https://github.com/wang-lizhi/QuantizationAwareDeepOptics.

6. Synthetic Simulation

To evaluate the effectiveness of our model, we conduct synthetic simulation on ICVL dataset [2]. We randomly divide the 201 spectral scenes into three parts, which are used for training (167 scenes), validation (17 scenes), and test (17 scenes), respectively. For each scene, we cut it into 9 overlapping patches with a size of 512×512. We set the

	PSNR↑			SSIM↑			RMSE↓			ERGAS↓		
Levels	DO	QDO	QDO+A	DO	QDO	QDO+A	DO	QDO	QDO+A	DO	QDO	QDO+A
Full	36.82	-	-	0.974	-	-	0.0167	-	-	10.35	-	-
32	36.75	36.80	36.82	0.974	0.974	0.974	0.0168	0.0167	0.0167	10.43	10.36	10.35
16	36.55	36.70	36.77	0.973	0.974	0.974	0.0172	0.0169	0.0168	10.69	10.50	10.42
8	35.17	36.21	36.44	0.969	0.972	0.973	0.0204	0.0180	0.0174	12.61	11.10	10.80
4	31.68	33.89	35.32	0.939	0.962	0.968	0.0319	0.0250	0.0205	19.16	14.85	12.35
2	28.56	33.42	33.97	0.891	0.952	0.958	0.0479	0.0288	0.0255	27.50	15.69	14.53

Table 1. The numerical comparison of the reconstruction results on the ICVL dataset among the conventional deep optics model (DO) that ignores DOE quantization in the fabrication [5, 15], the proposed QDO and QDO+A models. We compare the results with quantization levels from 2 to 32. Metrics of the ideal full-precision model are listed in the first row as references.

Levels	QDO ($\|H_q - H_f\|_1$)	QDO+A ($\|H_{aq} - H_f\|_1$)
32	20.707	0.763
16	44.075	1.965
8	94.192	1.822
4	202.510	18.911
2	285.000	188.010

Table 2. Quantization deviation (MAE, smaller is better) in nm between the full-precision height map and the quantized height map in proposed quantization-aware models with or without the adaptive mechanism.

Encoding	PSNR↑	SSIM↑	RMSE↓	ERGAS↓
Fresnel	27.41	0.869	0.0556	32.18
CASSI	30.66	0.899	0.0354	20.74
DO	31.68	0.939	0.0319	19.16
QDO	33.89	0.962	0.0250	14.85
QDO+A	35.32	0.968	0.0205	12.35

Table 3. Quantitative comparison among CASSI, Fresnel lens, the conventional deep optics model, the QDO model, and the QDO+A model.

physical parameters about the hardware system according to the physical experiment configurations in Sec. 7. Specifically, the DOE pixel size is $4\mu m \times 4\mu m$. The distance between the scene and the DOE is $1m$, and the distance between the DOE and the sensor plane is $50mm$. The DOE uses the refractive indices of SK1300 fused silica material, and the base plane thickness is set to $2mm$. We also measure the RGB response curve of the FLIR GS3-U3-41S4C-C camera and configure it as spectral response functions in the simulation.

6.1. Comparison with Conventional Deep Optics

Since the quantization deviation becomes insignificant when the quantization levels are more than 32, we perform synthetic simulations with from 2 to 32 quantization levels to compare the performance of different DO models on hyperspectral image reconstruction. For the conventional DO models in [5, 15], the DOE height map is optimized in full precision (32-bit float) in the training and quantized to multi-level style in the testing.

Tab. 1 shows the numerical results of the conventional model, the non-adaptive quantization-aware model and the adaptive quantization-aware model. The performance of the ideal model with full-precision DOE height map is also shown as references. We can see that both quantization-aware models outperform the conventional model at each quantization level, which demonstrates the effectiveness of the quantization-aware optimization for diffractive snapshot hyperspectral imaging. Further, the adaptive quantization-aware model performs much better than the non-adaptive quantization-aware model at each level, indicating the superiority of the adaptive mechanism.

Here, we show an in-depth analysis of the quantized DOE height maps H_q and H_{aq} in our models. We calculate the mean absolute error (MAE) between the quantized and the full precision height maps of the QDO and QDO+A models. The results in Tab. 2 illustrate that the adaptive mechanism effectively reduces the quantization deviation between the quantized height map and its full-precision weight, thus leading to a significant improvement on the hyperspectral image reconstruction to approximate the oracle performance.

Fig. 5 shows the reconstruction results of three representative images. To simultaneously present the results of all spectral bands, we convert the hyperspectral images to sRGB according to the spectral response function we use in simulation. Clearly, QDO and QDO+A can produce visually pleasant results with less artifact and sharper edges compared with the conventional model. We also plot spectral curves of random patches as shown in Fig. 5e. The spectral curves reconstructed by the proposed models are closer to the ground truth. Also, QDO+A has a distinct advantage over the QDO on both spatial image quality and spectral curve fidelity, which is consistent with the numerical comparison results.

(a) DO (b) QDO (c) QDO+A (d) Ground Truth (e) Spectral Curves

Figure 5. The visual comparison of the reconstruction results. Consistent with the numeric results in Tab. 1, visual perceptual results show that our quantization-aware models can effectively reduce the degradation of reconstruction results caused by DOE quantization, and the spectral results reconstructed by the QDO and the QDO+A model are also closer to the ground truth. Red rectangles in ground truth images indicate the area used for plotting the spectral curves.

6.2. Comparison with Other Optical Encodings

To evaluate the performance of our models in snapshot hyperspectral imaging, We compare with representative snapshot hyperspectral imaging system, *i.e.*, Fresnel lens and CASSI [3]. To make the comparison fair, we adopt the same reconstruction network as we use in Sec. 5 for all optical encoding models. Meanwhile, all models are trained for 50 epochs with the same optimizer configuration, and the quantization level number is set to 4 for the DO, QDO, and QDO+A.

Tab. 3 shows the reconstruction results of different encoding approaches. We can see that all three kinds of DO encodings have advantages over other snapshot hyperspectral encodings, and the proposed quantization-aware model and its adaptive mechanism further promote the advantages.

7. Physical Experiment

We fabricate the optimized DOE of the 4-level QDO+A model using multi-level lithography technologies. Owning to the quantization-aware model, the optimized DOE structure is already quantized, which means no extra quantization is needed and the height map can be directly used for DOE fabrication. Fig. 6 shows the detailed structure of the optimized DOE and its PSF in 31 spectral bands from $400nm$ to $700nm$. Then, we build the prototype hyperspectral imaging system using the fabricated DOE. The camera model and all physical parameters in fabrication experiments have been listed in the Sec. 6. The fabricated DOE is installed on the camera using a customized connector that supports C-mount on the camera side and holds the DOE $50mm$ away from the sensor plane. The diameter of the DOE is half an inch, and the diffractive pattern area size is $4.096mm$. The area without the diffractive pattern of the DOE is blocked by a customized steel aperture.

Fig. 7 shows the DOE we fabricated and the prototype camera we built. Fig. 8 shows the captured RGB images of the prototype camera and the reconstruction results. We also capture ground truth spectral curves with a commercial spectrometer (StallerNet Blue). We can see that our method can produce visually pleasant images and high-accuracy spectral curves.

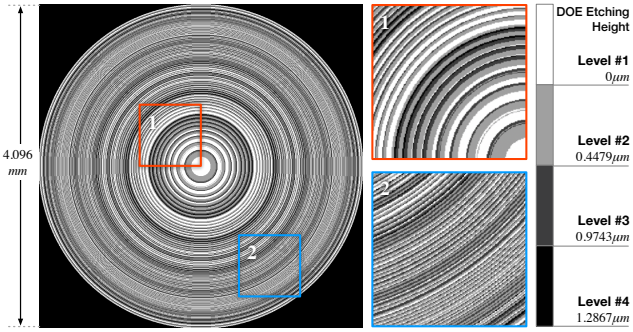

(a) The Optimized 4-level Height Map

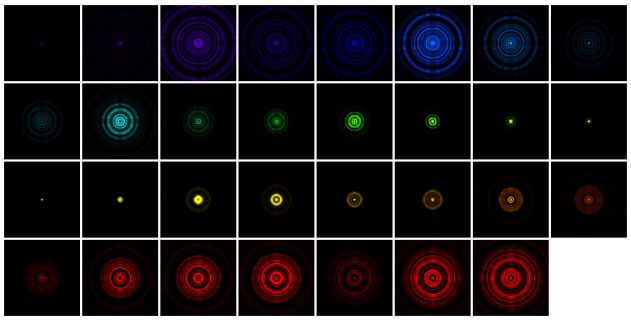

(b) The Corresponing PSF (from $400nm$ to $700nm$, 31 channels)

Figure 6. The height map optimized by the QDO+A 4-level and the corresponding PSF from 400nm to 700nm. This height map is naturally quantized and can be directly used for fabrication.

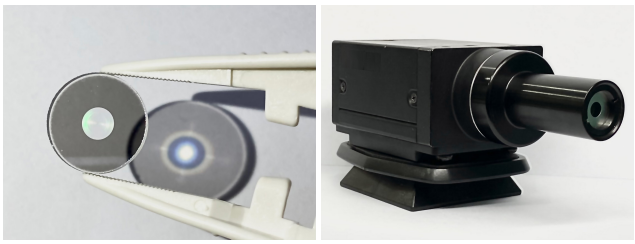

(a) Fabricated DOE (b) Prototype Camera

Figure 7. The fabricated DOE and the camera prototype we build for diffractive snapshot hyperspectral imaging.

(a) Captured RGB Images

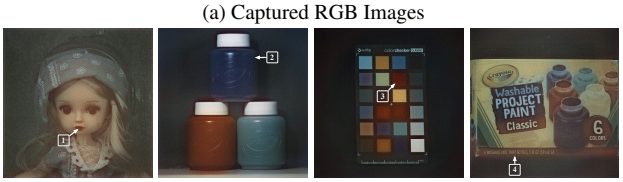

(b) Reconstruction (Visualized to RGB Image)

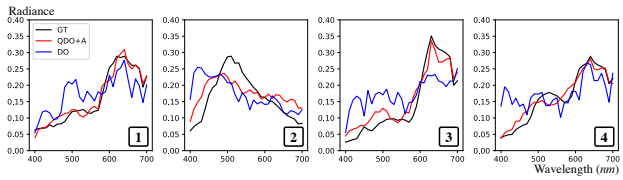

(c) Reconstruction Spectral Curves

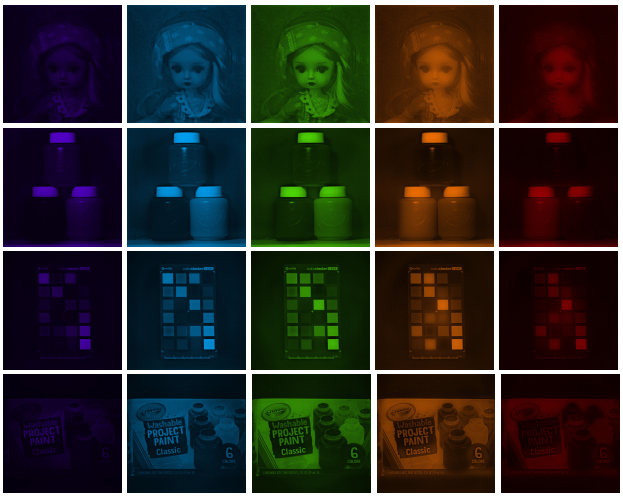

(d) Spectral Images ($420nm$, $470nm$, $530nm$, $620nm$, and $700nm$)

Figure 8. The captured encoded images, reconstruction results (RGB visualization), spectral curves, and spectral images in representative bands. Our system can produce visually pleasant hyperspectral images and accurate spectrum, indicating the effectiveness of our model in the physical system.

8. Discussion

Limitations. The proposed quantization-aware methods have limited effective scopes, and the approximation of the optical modeling might prevent the designed system from having a large valid field-of-view. Please check the supplementary material for a detailed discussion.

Conclusion. In this work, we proposed the quantization-aware deep optics for diffractive snapshot hyperspectral imaging, which can jointly optimizes a quantized DOE as the encoder and a reconstruction network as the decoder. We also presented an adaptive mechanism for this framework to further reduce the quantization deviation by adjusting the physical height of each level. We fabricated a prototype and built a real hardware systems according to the proposed models for diffractive snapshot hyperspectral imaging. Synthetic simulation and physical experiments have verified the effectiveness of our model. The quantization-aware deep optics are also applicable for other tasks, including low-level imaging and high-level vision. We would make more efforts on the generalization.

Acknowledgments

This work is supported by National Natural Science Foundation of China (62131003, 62072038, 61922014).

References

[1] Telmo Adão, Jonáš Hruška, Luís Pádua, José Bessa, Emanuel Peres, Raul Morais, and Joaquim Joao Sousa. Hyperspectral imaging: A review on uav-based sensors, data processing and applications for agriculture and forestry. *Remote Sensing*, 9(11):1110, 2017. 1

[2] Boaz Arad and Ohad Ben-Shahar. Sparse recovery of hyperspectral signal from natural rgb images. In *European Conference on Computer Vision*, pages 19–3corresponding author4, 2016. 5

[3] Gonzalo R Arce, David J Brady, Lawrence Carin, Henry Arguello, and David S Kittle. Compressive coded aperture spectral imaging: An introduction. *IEEE Signal Processing Magazine*, 31(1):105–115, 2013. 1, 2, 7

[4] V Backman, Michael B Wallace, LT Perelman, JT Arendt, R Gurjar, MG Müller, Q Zhang, G Zonios, E Kline, T McGillican, et al. Detection of preinvasive cancer cells. *Nature*, 406(6791):35–36, 2000. 1

[5] Seung-Hwan Baek, Hayato Ikoma, Daniel S Jeon, Yuqi Li, Wolfgang Heidrich, Gordon Wetzstein, and Min H Kim. Single-shot hyperspectral-depth imaging with learned diffractive optics. In *International Conference on Computer Vision*, pages 2651–2660, 2021. 1, 2, 4, 6

[6] Seung-Hwan Baek, Incheol Kim, Diego Gutierrez, and Min H Kim. Compact single-shot hyperspectral imaging using a prism. *ACM Transactions on on Graphics*, 36(6):217, 2017. 1

[7] Yu Bai, Yu-Xiang Wang, and Edo Liberty. Proxquant: Quantized neural networks via proximal operators. In *International Conference on Learning Representations*, pages 1–17, 2019. 3

[8] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 2

[9] Marcus Borengasser, William S Hungate, and Russell Watkins. *Hyperspectral remote sensing: principles and applications*. CRC press, 2007. 1

[10] Xun Cao, Hao Du, Xin Tong, Qionghai Dai, and Stephen Lin. A prism-mask system for multispectral video acquisition. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 33(12):2423–2435, 2011. 1

[11] Xun Cao, Tao Yue, Xing Lin, Stephen Lin, Xin Yuan, Qionghai Dai, Lawrence Carin, and David J Brady. Computational snapshot multispectral cameras: toward dynamic capture of the spectral world. *IEEE Signal Processing Magazine*, 33(5):95–108, 2016. 1

[12] Julie Chang and Gordon Wetzstein. Deep optics for monocular depth estimation and 3d object detection. In *International Conference on Computer Vision*, pages 10193–10202, 2019. 1, 2

[13] Tenn F Chen, Gladimir VG Baranoski, Bradley W Kimmel, and Erik Miranda. Hyperspectral modeling of skin appearance. *ACM Transactions on on Graphics*, 34(3):1–14, 2015. 1

[14] David L Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006. 2

[15] Xiong Dun, Hayato Ikoma, Gordon Wetzstein, Zhanshan Wang, Xinbin Cheng, and Yifan Peng. Learned rotationally symmetric diffractive achromat for full-spectrum computational imaging. *Optica*, 7(8):913–922, 2020. 1, 2, 3, 4, 6

[16] Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. A survey of quantization methods for efficient neural network inference. *arXiv preprint arXiv:2103.13630*, 2021. 2

[17] Joseph W Goodman. Introduction to fourier optics. 3rd. *Roberts and Company Publishers*, 2005. 3

[18] Felix Heide, Qiang Fu, Yifan Peng, and Wolfgang Heidrich. Encoded diffractive optics for full-spectrum computational imaging. *Scientific Reports*, 6(1):1–10, 2016. 1

[19] Tao Huang, Weisheng Dong, Xin Yuan, Jinjian Wu, and Guangming Shi. Deep gaussian scale mixture prior for spectral compressive imaging. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 16216–16225, 2021. 2

[20] Hayato Ikoma, Cindy M Nguyen, Christopher A Metzler, Yifan Peng, and Gordon Wetzstein. Depth from defocus with learned optics for imaging and occlusion-aware depth estimation. In *IEEE International Conference on Computational Photography*, pages 1–12, 2021. 1, 2

[21] John James. *Spectrograph design fundamentals*. Cambridge University Press, 2007. 1

[22] Daniel S Jeon, Seung-Hwan Baek, Shinyoung Yi, Qiang Fu, Xiong Dun, Wolfgang Heidrich, and Min H Kim. Compact snapshot hyperspectral imaging with diffracted rotation. *ACM Transactions on on Graphics*, 38(4):1–13, 2019. 1, 2

[23] Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018. 2

[24] Cong Leng, Zesheng Dou, Hao Li, Shenghuo Zhu, and Rong Jin. Extremely low bit neural network: Squeeze the last bit out with admm. In *Association for the Advancement of Artificial Intelligence*, 2018. 3

[25] Xing Lin, Yebin Liu, Jiamin Wu, and Qionghai Dai. Spatial-spectral encoded compressive hyperspectral imaging. *ACM Transactions on on Graphics*, 33(6):233, 2014. 1

[26] Zhigang Liu and Matthew Mattina. Learning Low-precision Neural Networks without Straight-Through Estimator (STE). In *International Joint Conferences on Artificial Intelligence*, pages 3066–3072, 2019. 3, 4

[27] Guolan Lu and Baowei Fei. Medical hyperspectral imaging: a review. *Journal of Biomedical Optics*, 19(1):010901, 2014. 1

[28] Chenguang Ma, Xun Cao, Xin Tong, Qionghai Dai, and Stephen Lin. Acquisition of high spatial and spectral resolution video with a hybrid camera system. *International Journal of Computer Vision*, 110(2):141–155, 2014. 1

[29] Christopher A Metzler, Hayato Ikoma, Yifan Peng, and Gordon Wetzstein. Deep optics for single-shot high-dynamic-range imaging. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1375–1385, 2020. 1, 2

[30] Brian Morgan, Christopher M Waits, John Krizmanic, and Reza Ghodssi. Development of a deep silicon phase fresnel

lens using gray-scale lithography and deep reactive ion etching. *Journal of microelectromechanical systems*, 13(1):113–120, 2004. 2

[31] Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. Up or down? adaptive rounding for post-training quantization. In *International Conference on Machine Learning*, pages 7197–7206. PMLR, 2020. 3

[32] Zhihong Pan, Glenn Healey, Manish Prasad, and Bruce Tromberg. Face recognition in hyperspectral images. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 25(12):1552–1560, 2003. 1

[33] Yifan Peng, Qiang Fu, Felix Heide, and Wolfgang Heidrich. The diffractive achromat full spectrum computational imaging with diffractive optics. *ACM Transactions on on Graphics*, 35(4):31:1–31:11, 2016. 1

[34] Antonio Polino, Razvan Pascanu, and Dan Alistarh. Model compression via distillation and quantization. *arXiv preprint arXiv:1802.05668*, 2018. 2

[35] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 5

[36] Vincent Sitzmann, Steven Diamond, Yifan Peng, Xiong Dun, Stephen Boyd, Wolfgang Heidrich, Felix Heide, and Gordon Wetzstein. End-to-end optimization of optics and image processing for achromatic extended depth of field and super-resolution imaging. *ACM Transactions on on Graphics*, 37(4):1–13, 2018. 1, 2

[37] Qilin Sun, Ethan Tseng, Qiang Fu, Wolfgang Heidrich, and Felix Heide. Learning rank-1 diffractive optics for single-shot high dynamic range imaging. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1386–1396, 2020. 1, 2, 4

[38] Shiyu Tan, Yicheng Wu, Shoou-I Yu, and Ashok Veeraraghavan. Codedstereo: Learned phase masks for large depth-of-field stereo. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7170–7179, 2021. 2

[39] Ethan Tseng, Shane Colburn, James Whitehead, Luocheng Huang, Seung-Hwan Baek, Arka Majumdar, and Felix Heide. Neural nano-optics for high-quality thin lens imaging. *Nature communications*, 12(1):1–7, 2021. 2

[40] Ethan Tseng, Ali Mosleh, Fahim Mannan, Karl St-Arnaud, Avinash Sharma, Yifan Peng, Alexander Braun, Derek Nowrouzezahrai, Jean-François Lalonde, and Felix Heide. Differentiable compound optics and processing pipeline optimization for end-to-end camera design. *ACM Transactions on on Graphics*, 40(2):1–19, 2021. 2

[41] Ashwin Wagadarikar, Renu John, Rebecca Willett, and David Brady. Single disperser design for coded aperture snapshot spectral imaging. *OSA Applied Optics*, 47(10):B44–B51, 2008. 1

[42] Lizhi Wang, Chen Sun, Ying Fu, Min H Kim, and Hua Huang. Hyperspectral image reconstruction using a deep spatial-spectral prior. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8032–8041, 2019. 2

[43] Lizhi Wang, Chen Sun, Maoqing Zhang, Ying Fu, and Hua Huang. Dnu: Deep non-local unrolling for computational spectral imaging. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1661–1671, 2020. 2

[44] Lizhi Wang, Zhiwei Xiong, Hua Huang, Guangming Shi, Feng Wu, and Wenjun Zeng. High-speed hyperspectral video acquisition by combining nyquist and compressive sampling. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 41(4):857–870, 2018. 2

[45] Lizhi Wang, Shipeng Zhang, and Hua Huang. Adaptive dimension-discriminative low-rank tensor recovery for computational hyperspectral imaging. *International Journal of Computer Vision*, 129(10):2907–2926, 2021. 2

[46] Fengchao Xiong, Jun Zhou, and Yuntao Qian. Material based object tracking in hyperspectral videos. *IEEE Transactions on Image Processing*, 29:3719–3733, 2020. 1

[47] Penghang Yin, Jiancheng Lyu, Shuai Zhang, Stanley Osher, Yingyong Qi, and Jack Xin. Understanding straight-through estimator in training activation quantized neural nets. *arXiv preprint arXiv:1903.05662*, 2019. 3

[48] Xin Yuan, David J Brady, and Aggelos K Katsaggelos. Snapshot compressive imaging: Theory, algorithms, and applications. *IEEE Signal Processing Magazine*, 38(2):65–88, 2021. 2

[49] Zhengxin Zhang, Qingjie Liu, and Yunhong Wang. Road extraction by deep residual u-net. *IEEE Geoscience and Remote Sensing Letters*, 15(5):749–753, 2018. 5

[50] Bohan Zhuang, Chunhua Shen, Mingkui Tan, Lingqiao Liu, and Ian Reid. Towards effective low-bitwidth convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7920–7928, 2018. 2