Revisiting Light Field Rendering with Deep Anti-Aliasing Neural Network

Gaochang Wu, Yebin Liu, Member, IEEE, Lu Fang, Senior Member, IEEE, and Tianyou Chai, Fellow, IEEE

Abstract—The light field (LF) reconstruction is mainly confronted with two challenges, large disparity and non-Lambertian effect. Typical approaches either address the large disparity challenge using depth estimation followed by view synthesis or eschew explicit depth information to enable non-Lambertian rendering, but rarely solve both challenges in a unified framework. In this paper, we revisit the classic LF rendering framework to address both challenges by incorporating it with advanced deep learning techniques. First, we analytically show that the essential issue behind the large disparity and non-Lambertian challenges is the aliasing problem. Classic LF rendering approaches typically mitigate the aliasing with a reconstruction filter in the Fourier domain, which is, however, intractable to implement within a deep learning pipeline. Instead, we introduce an alternative framework to perform anti-aliasing reconstruction in the image domain and prove in theory the comparable and even superior efficacy on the aliasing issue. To explore the full potential, we then embed the anti-aliasing framework into a deep neural network through the design of an integrated architecture and trainable parameters. The network is trained through end-to-end optimization using a peculiar training set, including regular LFs and unstructured LFs. The proposed deep learning pipeline shows a substantial superiority on solving both the large disparity and the non-Lambertian challenges compared with other state-of-the-art approaches. In addition to the view interpolation for a LF, we also show that the proposed pipeline also benefits light field view extrapolation.

Index Terms—Light field reconstruction, light field rendering, deep learning, view extrapolation.

1 INTRODUCTION

As an alternative to traditional 3D scene representation using scene geometry (or depth) and texture (or reflectance), light field (LF) achieves photorealistic view synthesis in real-time using LF rendering technology [1], [2]. This high-quality rendering requires the disparities between adjacent views to be less than one pixel, i.e., the so-called densely-sampled LF. Unfortunately, the practical scenarios including dynamic scene [3] or limited acquisition time [4] impose insufficient sampling in the angular dimension. The quality of the rendered novel views is inevitably perturbed by the large disparity (range) in the sampled LF. In addition, the potential non-Lambertian effect in the scene, such as jewellery, fur, glass and face, will further aggravate this side effect [5], [6].

To overcome these two technical challenges, i.e., large disparity and non-Lambertian effect, classic LF rendering approaches [5], [6], [7] tended to treat the plenoptic sampling and reconstruction in a signal processing way. They may design a powerful anti-aliasing filter to reconstruct the LF in the Fourier domain, as illustrated in Fig. 2(b). However, it relies on a carefully designed filter determined by the depth information and the aliasing degree, which are prone to blurry results due to the high-frequency loss [8]. With the boost of tremendous deep learning techniques, state-of-the-art approaches achieve more promising result,

- Gaochang Wu and Tianyou Chai are with the State Key Laboratory of Synthetical Automation for Process Industries, Northeastern University, Shenyang 110819, P. R. China. E-mail: ahwgc2009@163.com, tychai@mail.neu.edu.cn.
- Yebin Liu is with Department of Automation, Tsinghua University, Beijing 100084, P. R. China. E-mail: liuyebin@mail.tsinghua.edu.cn.
- Fang Lu is with Tsinghua-Berkeley Shenzhen Institute, Shenzhen 518055, P. R. China. E-mail: fanglu@sz.tsinghua.edu.cn.

Corresponding Author: Yebin Liu, Fang Lu.



Fig. 1. Schematic diagram of the proposed Deep Anti-Aliasing Network (DA^2N) with embedding shearing, downscaling and prefiltering. In the result, we use input views indicated by the red lines for the EPI reconstruction, including interpolation and extrapolation.

either depth estimation followed by view synthesis [9], [10] or direct LF reconstruction [11], [12], [13]. However, existing view synthesis or LF reconstruction methods can hardly solve the aforementioned challenges together. For instance, Kalantari *et al.* [10] developed sequential networks to infer the scene depth and synthesize novel view via the four corner sub-aperture images. Since the depth estimation is usually based on the Lambertian assumption, it will lead to ambiguous result in a non-Lambertian surface (see Fig. 10). Alternatively, some other approaches [12], [13] eschew explicit depth information to enable non-Lambertian rendering, while they suffer from either aliasing or blurring effect under the large disparity (see Fig. 9 and Fig. 10). It is apparent that the effect caused by either the large disparity or the non-Lambertian effect is ubiquitous in real-world, neglecting either of them will certainly lead to quality degeneration in the reconstructed LF.

In this paper, we consider the two challenges, large disparity and non-Lambertian effect, in a unified deep learning framework. The detailed analysis in the Fourier domain reveals that the essential issue behind the two challenge is the aliasing problem [5], [6]. The intuitive way of mimicking the anti-aliasing filter [7], [8] via a neural network in the Fourier domain is intractable due to the requirement of a global spectrogram perception (Sec. 3.1). Our insight is decomposing the classical anti-aliasing filtering in the Fourier domain into a sequence of operations, including "shearing, downscaling and prefiltering" in the image domain, and formulate them with the deep learning pipeline. Compared with other available learning-based methods that resort to a deeper representation, our shearing-downscalingprefiltering framework owns the advantage of explicitly and efficiently aliasing handling from the essential structure design, and thus, is more competitive in solving large disparity and non-lambertian scenarios.

Specifically, the shear operation [14], [15] converts the anti-aliasing filter with a complicated representation in the Fourier domain into a simple low-pass filter (i.e., the Gaussian filter) in the image domain via plane (sub-aperture image) sweep. Subsequently, rather than directly applying the Gaussian filter, we address the aliasing issue by combining downscaling and prefiltering in the image domain. The downscaling operation indirectly reduces the disparity range [5] by resizing the spatial resolutions of the input LF. And the prefiltering operation further suppresses the aliasing high-frequencies introduced by the non-Lambertian effect. The proposed shearing-downscalingprefiltering framework is theoretically proved to own a higher performance than the conventional reconstruction filter in the Fourier domain.

With the guidance of the sequential operations, the designed network is able to explicitly and efficiently solve both challenges while also reasoning in an end-to-end manner. We achieve this through the following features. First, we apply a multi-branch structure to implement shear operation with different candidates. Accordingly, a fusion network is employed to post-fuse the reconstructed LFs with different shear amounts and produce the high angular resolution LF without depth estimation. Second, we propose a trainable Laplacian decomposition to formulate the downscaling operation and keep the high-frequencies simultaneously. Third, we use a convolutional layer in the spatial dimension to perform prefilter operation, which learns to handle the aliasing high-frequencies from training data. The overall network is shown in Fig. 1, and is termed as Deep Anti-Aliasing Network, or DA²N for short.

To further boost the performance, we begin by training the DA²N from scratch on regular LFs and then fine-tune the network using pseudo epipolar plane images (EPIs) from unstructured LFs (Sec. 5.1). On the one hand, we indicate that pseudo EPIs and regular EPIs with the non-Lambertian effect not just have the same visual features, but more importantly, share the same properties regarding to the aliasing issue. On the other hand, unstructured LF is an easily accessible source (e.g., by a hand-held camera) and is able to offer a large disparity compared with a regular LF (e.g., Lytro LF and HCI synthetic LF [16]). The technical contributions are summarized as follows:

- We propose a novel end-to-end Deep Anti-aliasing Neural Network (DA²N) to solve the two challenges, large disparity and non-Lambertian effect, in light field rendering. For the network training, the pseudo EPIs from unstructured LFs are particularly used.
- We provide a theoretical analysis of the aliasing problem that leads to the two challenges. It can be proved that the proposed "shearing, downscaling and prefiltering" in image domain owns a higher performance than the conventional reconstruction filter in the Fourier domain.
- We demonstrate the superiority of the DA²N for high-quality LF reconstruction with both non-Lambertian effect and large disparity (up to 14 pixels), as shown in Fig. 8. In addition, by combining with a discriminator, DA²N can be further extended to view extrapolation, as shown in Fig. 1.

2 RELATED WORK

Depth-based view synthesis. These kind of approaches typically first estimate the depth of a scene [14], [17], [18], [19], such as structure tensor-based local direction estimation in the epipolar plane image (EPI) domain [20], and phase-based depth estimation in the sub-aperture image domain [21]. Then the input views are warped to novel viewpoints and blended in different manners, e.g., soft blending [22] and learning-based synthesis [23]. In recent years, some deep learning based methods for directly maximizing the quality of synthetic views have been presented. Flynn et al. [9] proposed a deep learning method to synthesize novel views using a sequence of images with wide baselines. Kalantari et al. [10] used two sequential convolutional neural networks to model depth and color estimation simultaneously by minimizing the error between synthetic views and ground truth images. Zhou et al. [24] trained a network that infers alpha and multiplane images. The novel views are synthesized using homography and alpha composition. However, all these approaches are based on the Lambertian assumption without explicitly addressing the non-Lambertian challenge.

Reconstruction without explicit depth. The LF reconstruction without using explicit depth information or scene geometry is based on sampling and consecutive reconstruction of the plenoptic function [25], [26], [27]. Typical approaches focus on an anti-aliasing reconstruction in the Fourier domain using signal processing techniques, e.g., quadra-linear interpolation [1] and depth corrected interpolation [2]. An initial work by Chai et al. [5] derived a plenoptic sampling theory, which indicates that the recovering of LF signal can be realized by applying an anti-aliasing reconstruction filter bounded within a double-wedge shape (please refer to Fig. 2(b)). Zhang and Chen [6] further showed that the non-Lambertian effect in the acquired LF will introduce spectra shifting along the angular dimension, and thus, an expansionary reconstruction filter is required for the LF reconstruction. Instead of using a fixed filter pattern, Vagharshakyan et al. [7] proposed to adopt a composited reconstruction filter packet to remove the high-frequency spectrums that introduce aliasing effects in Fourier domain.



Fig. 2. Fourier analysis of shear, downscaling (spatial) and prefilter operations. The EPI for demonstration is composed of three Lambertian points (A, C and D) and one non-Lambertian point (B) with different disparities. (a) Original EPI (top), its Fourier transform after taking the logarithm (middle) and diagram (down). The non-Lambertian point will introduce expansionary spectra on each side of the Ω_s axis [6]; (b) Spectrum replicas appear when the EPI is angularly down-sampled. The overlapped spectrum and its replicas will lead to aliasing effect. An ideal reconstruction filter (parallelogram) can be obtained using the optimal rendering disparity (depth) shown by the dashed line; (c) Shearing the EPI with a proper value forces the optimal rendering disparity of the reconstruction filter to rotate towards the Ω_u axis; (d) With a downscaling (spatial) operation, the original spectrum will shrink towards the Ω_u axis. And the separated spectra substantially mitigates the aliasing issue; (e) We show an EPI with directly up-sampling (middle), see the close-up for the detail. We can effectively decrease the shape parameter σ of the prefilter by using shear and downscaling operations (bottom).

For explicit aliasing handling, Levoy and Hanrahan [1] propose to use an optical prefiltering method that acts as a 4D low-pass filter. Alternatively, Chai *et al.* [5] and Stewart *et al.* [8] further indicated that a computational prefiltering method can be achieved by using a discrete low-pass filter after the LF sampling.

Based on the recent deep learning techniques, some depth-independent reconstruction approaches were proposed [11], [12], [13], [28]. Yeung et al. [13] applied a coarseto-fine model using a 4D convolutional network for LF reconstruction followed by an encoder network for view refinement. In spite of the large receptive field in the network, the reconstructed LF can also appear aliasing effects in large disparity cases (as shown in Fig. 9). To explicitly addressing the aliasing effects, Wu et al. [12] proposed to apply a prefilter operation before feeding the EPI into the network. However, limited by the postprocessing using non-blind deblur, the approach is prone to reconstruct a LF with blurring effect (see Fig. 7). By contrast, we model the prefilter operation within the neural network by adopting 1D convolutional layers with different kernel size, and design a trainable Laplacian decomposition for the balance between aliasing handling and high-frequency maintaining (please refer to Sec. 4.1 for detail).

3 FOURIER ANALYSIS: THE ALIASING PROBLEM

In this section, we present a Fourier analysis on the large disparity and non-Lambertian challenges based on the plenoptic theory [5], [6]. It reveals that the aliasing problem is the essential issue behind those challenges (Sec. 3.1), where the aliasing degree is determined by the highest frequency of the non-Lambertian effect and the largest disparity. We further introduce an effective framework containing shearing, downscaling and prefiltering operations for antialiasing LF reconstruction (Sec. 3.2). Here, each operation is accomplished in the image domain, and thus, can be efficiently formulated with a learning-based pipeline.

3.1 The large disparity and non-Lambertian challenges

For a 4D LF with two spatial dimensions (u, v) and two angular dimensions (s, t) represented as L(u, v, s, t), an EPI can be acquired by gathering horizontal lines with fixed v^* along a constant camera coordinate t^* , denoted as $E_{v^*,t^*}(u, s)$ (or $E_{u^*,s^*}(v, t)$ similarly). For simplicity, we denote the extracted EPI as E_{HR} , where HR stands for high angular resolution. The low angular resolution EPI E_{LR} can be considered as its undersampled version, i.e., $E_{LR} = E_{HR} \downarrow$, where \downarrow denotes the downsampling operation in the angular dimension (*s* or *t*).

Consider a simple scene composed of three Lambertian points (A, C and D) and a non-Lambertian point (B) located at different depths. Fig. 2(a) (top) shows an EPI extracted from a LF of the scene, where the disparities of point A, Band C are within 1 pixel range while D is larger than 1 pixel. The Fourier transform and its simplified diagram are shown in the middle and bottom of Fig. 2(a), respectively. The spectral support of the corresponding line with disparity d is defined by $d\Omega_u + \Omega_s = 0$, where Ω_u and Ω_s denote the frequencies along the u and s dimensions. Zhang and Chen [6] concluded that a non-Lambertian point will introduce expansionary spectra on each side of the Ω_s axis,

$$d\Omega_u + \Omega_s = \pm \beta / Z,\tag{1}$$

where β is the band width of the band-limited signal caused by the non-Lambertian effect [6], and *Z* is the depth of that non-Lambertian point. The region with red in Fig. 2(a) (bottom) demonstrates the non-Lambertian effect in the Fourier domain.

An initial work on LF sampling by Chai *et al.* [5] inferred that the down-sampling in the angular dimension will lead to spectrum replicas. Namely, the overlapping of the original spectrum and its replicas will lead to aliasing effect, as shown in Fig. 2(b). As indicated in [6] and also by Eqn. 1, disparity d and the highest frequency of non-Lambertian effect β are two independent factors regarding to the aliasing

| Main symbol used in the paper. | | | | | | |
|--------------------------------|--|------------|--------------------|--|--|--|
| Symbol | Definition | Symbol | Definition | | | |
| L | 4D Light field | E | 2D EPI | | | |
| u, v | Spatial dimension | s,t | Angular dimension | | | |
| d | Disparity | Z | Depth | | | |
| α_h | Shear amount | α_u | Downscaling factor | | | |
| ${\cal F}$ | Fourier transform | γ | Spectrum amplitude | | | |
| Symbol | Definition | | | | | |
| β | Band width of the signal caused by non-Lambertian effect | | | | | |
| α_s | Upsampling factor in the angular dimension | | | | | |
| f | Functions defined in the image domain | | | | | |
| ϕ | Layer in a neural network | | | | | |
| θ | Trainable parameters in a neural network | | | | | |

issue. In other words, the non-Lambertian effect will also introduce aliasing even when the corresponding disparity is less than 1-pixel, and the larger disparity or the higher band width of the band-limited signal (caused by the non-Lambertian effect) may worsen the aliasing problem. Therefore, the key to solve the large disparity and non-Lambertian challenges in a unified framework is the aliasing handling.

To address the aliasing issue, a reconstruction filter can be designed for aliasing-free reconstruction of the high angular resolution EPI, as shown by the parallelogram in Fig. 2(b). The reconstruction filter is defined by the optimal rendering disparity $d_{opt} = \frac{d_{min}+d_{max}}{2}$, and the cut-off frequency is determined by the aliasing degree, where d_{min} and d_{max} are the minimum and maximum disparities. By incorporating the reconstruction filter with deep learning techniques in the Fourier domain, one may design a network to determine the optimal disparity and the cut-off frequency. However, such a procedure requires a global perception of the entire spectrogram, i.e., the receptive field [29] of the network should cover the entire input image, which inevitably introduces a mass of parameters.

3.2 Framework for anti-aliasing reconstruction

Instead of adopting a reconstruction filter in the *Fourier domain* as used in [5], [6], [7], which is intractable to implement within a deep learning pipeline, we derive a novel antialiasing framework composed of three sequential operations in the *image domain*, including shearing, downscaling and prefiltering. The proposed framework shows high efficiency in solving the large disparity and non-Lambertian challenges.

Shearing operation. It translates the disparity range of the input EPI to be centered at zero, so as to minimize the maximum (absolute) disparity value [14], [15]. The shear operation is achieved by shifting each sub-aperture view according to the following equation

$$f_h(E(u,s);\alpha_h) = E(u+s \cdot \alpha_h,s), \tag{2}$$

where f_h is the shearing operation and α_h is the shear amount. By shearing the EPI with the optimal rendering disparity d_{opt} , the minimum and maximum disparities become $d'_{\min} = \frac{d_{\min} - d_{\max}}{2}$, $d'_{\max} = \frac{d_{\max} - d_{\min}}{2}$, and the optimal rendering disparity $d'_{opt} = 0$. Fig. 2(c) illustrates the EPI, the Fourier transform and the simplified diagram after the shearing operation.

The shearing operation is a key element of the framework, due to the facts that: 1) it is highly effective for a local EPI patch with large disparity yet small disparity range; 2) it ensures the reconstruction filter to be easily implemented in the image domain, e.g., a Gaussian function. Note that although the shearing may reduce the maximum disparity to a certain extent, the size of disparity range remains the same, which still induces the aliasing problem.

Downscaling operation. Instead of simply adopting a Gaussian filter for anti-aliasing, we particularly apply the downscaling operation to resize the EPI to a reduced spatial size (pixel resolution). Note that it is a different operation compared with the angular downsampling, we, therefore, use the term "downscaling" for distinction. The downscaling operator shows superiority comparing against the conventional reconstruction filter, whose proof is presented in the supplementary. Along with the spatial downscaling, the disparity range is reduced. Therefore, by rescaling a sheared EPI with factor α_u , the spectral support will be

$$d'/\alpha_u \Omega_u + \Omega_s = \pm \beta/Z. \tag{3}$$

We have $\beta = 0$ and $\beta > 0$ for Lambertian and non-Lambertian cases, respectively. And the minimum and maximum disparities are $d''_{\min} = \frac{d_{\min} - d_{\max}}{2\alpha_u}$, $d''_{\max} = \frac{d_{\max} - d_{\min}}{2\alpha_u}$. The downscaling operation forces the original spectrum

The downscaling operation forces the original spectrum to shrink towards the Ω_u axis, separating the original spectrum from its replicas to mitigate the aliasing problem, as shown in Fig. 2(d). Afterwards, the width of the reconstruction filter will be increased by a factor of α_u , as shown in Fig. 2(d) (middle and bottom). One may notice that the pure downscaling operation will cause high-frequency loss according to the downscaling factor. Fortunately, by taking advantage of the deep learning techniques, we can adopt a multi-scale strategy for the balance between aliasing handling and high-frequency maintaining.

Prefiltering operation. The prefiltering operation alleviates the aliasing introduced by the reduced disparity range and the non-Lambertian effect via suppressing the highfrequency components in an EPI [30]. In the following, we empirically show that the combination of the downscaling and prefiltering is more effective than using a prefilter merely or the conventional reconstruction filter [5], [6].

In this paper, we implement the prefiltering using a Gaussian kernel κ in image domain, which is equivalent to the reconstruction filtering after taking the shearing with optimal disparity [5]. Note that the Gaussian kernel κ in image domain can be transformed as \mathcal{F}_{κ} , a Gaussian function in Fourier domain

$$\kappa(u;\sigma) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{u^2}{2\sigma^2}} \Leftrightarrow \mathcal{F}_{\kappa}(\Omega_u;\sigma) = e^{-\frac{\Omega_u^2}{1/(2\pi^2\sigma^2)}}, \quad (4)$$

where σ is the shape parameter. Then the prefiltering or reconstruction filtering in Fourier domain equals to the convolution with kernel in image domain.

The prefilter \mathcal{F}_{κ} is designed via the criteria that, the amplitude of a reference aliasing point P_a should not exceed a baseline spectrum amplitude γ after being filtered

$$\mathcal{F}_{E_{LR}}(P_a)\mathcal{F}_{\kappa}(\Omega_{u,P_a};\sigma) \le \gamma,\tag{5}$$

where $\mathcal{F}_{E_{LR}}$ is the Fourier transform of the EPI E_{LR} , the reference aliasing P_a has the lowest spatial frequency, Ω_{u,P_a} is the Ω_u coordinate of point P_a , as shown in Fig. 2(c) and



Fig. 3. Architecture of the proposed deep anti-aliasing network (DA²N) based on shearing, downscaling and prefiltering. The DA²N contains a reconstruction module for high angular resolution EPI and a fusion module for blending the EPIs with different shear amounts.

(d). We can set different γ for different degrees of antialiasing. Substitute Eqn. 4 into Eqn. 5, we then have the shape parameter σ as

$$\sigma \ge \sqrt{\frac{1}{2\pi^2 \cdot \Omega_{u,P_a}^2} \ln\left(\frac{\mathcal{F}_{E_{LR}}(P_a)}{\gamma}\right)},\tag{6}$$

where $\Omega_{u,P_a} = \alpha_u [(\Omega_{s,P_a} \pm \beta_{P_a}/Z_{P_a})/(-d'_{P_a})]$ according to Eqn. 3. As the depth will not be influenced by the shearing nor the downscaling operation, we consider β_{P_a}/Z_{P_a} as a constant. Also, the vertical coordinate Ω_{s,P_a} and the disparity after the shearing d'_{P_a} are fixed. We then have

$$\sigma \propto \frac{1}{\alpha_u} \sqrt{\ln\left(\frac{\mathcal{F}_{E_{LR}}(P_a)}{\gamma}\right)}$$

The $\sigma - \alpha_u$ curves with different baseline amplitude γ in range [5, 25] are plotted in the bottom of Fig. 2(e). On the one hand, the shape parameter σ will be smaller with the increasing of the downscaling factor α_u (as shown in the dotted line). On the other hand, the amplitude of point P_a will also be attenuated as the increase of α_u . Therefore, the final σ of the prefilter will be much smaller after the downscaling operation. In other words, the combination of downscaling and per-filtering is more effective than simply applying a prefilter or a reconstruction filter on the original signal. The top of Fig. 2(e) shows the reconstructed EPI using the introduced framework including shear, downscaling and per-filter operations.

4 DEEP ANTI-ALIASING NEURAL NETWORK

In this section, we present DA²N, a Deep Anti-Aliasing Neural Network based on the Fourier analysis on the shearing, downscaling and prefiltering in Sec. 3. The network contains a reconstruction module for high angular resolution EPI based on the shearing, downscaling and prefiltering, and a fusion module for blending the EPIs (feature maps) with different shear amounts. Fig. 3 visualizes the network structure of DA²N, which infers high angular resolution EPIs in an end-to-end manner. For reconstructing a 4D LF from the 2D EPIs, a two-step strategy simplified from the hierarchical reconstruction strategy in [12] is adopted, as illustrated in Fig. 4. EPIs $E_{v^*,t^*}(u,s)$ extracted from the horizontal views (marked as blue in Fig. 4(b)) are applied for reconstructing LF with high resolution in the *s* dimension. Similarly, EPIs $E_{u^*,s^*}(v,t)$ extracted from the vertical views (marked as green in Fig. 4(c)) are applied for reconstructing LF with high resolution in the *t* dimension.



Fig. 4. Two-step strategy for reconstructing a 4D LF from 2D EPIs.

4.1 Reconstruction module

The reconstruction module embeds the shearing, downscaling and prefiltering operations within a deep learning framework. The first step of the module is shearing the input low angular resolution EPI E_{LR} with a set of shear amounts H, generating multiple streams $f_h(E_{LR}; \alpha_h)$ as shown in Fig. 3. Next, the sheared inputs $f_h(E_{LR}; \alpha_h)$ are fed into a reconstruction net (Fig. 3, top right) to infer high angular resolution feature maps based on the downscaling and prefiltering. Note that the weights in the reconstruction net are shared for the input with different shear amounts. The final step of the reconstruction module is recovering the high angular resolution feature maps from the former shear operation. We implement this step through another shear operation with the amount $-\alpha_h/\alpha_s$, where α_s is the upsampling factor in the angular dimension.

To construct a **shear layer** in the neural network, we extend the operation described in Eqn. 2 to a general form

$$f_h(\phi(u, s, c); \alpha_h) = \phi(u + (s - \frac{S}{2}) \cdot \alpha_h, s, c),$$

where *S* is the number of pixels in the angular dimension, $\alpha_h \in H$ is the shear amount and $\phi(u, s, c)$ is a 3D tensor with a spatial dimension *u*, an angular dimension *s* and a channel dimension *c*. Note that ϕ can be the input EPI as well as a feature in the neural network. We use bilinear interpolation for the shear layer described in this paper. And the blank elements in a sheared tensor are set to zero. With the proposed shear layer, the network is able to reconstruct LF with view consistency when encounter with the large disparity challenge, see the comparison demonstrated in Fig. 5 (b) and (d).

A similar utilization of the shear operation in a learningbased method can be found in [31]. However, the proposed pipeline is trained to directly reconstruct a high angular resolution EPI, but the network in [31] is designed to decide whether the EPI is sheared by the optimal rendering disparity. In addition, the proposed pipeline has two fundamental differences in terms of the shear layer as follows: 1) The shear layer in the proposed pipeline is applicable not only



Fig. 5. The reconstructed LF (or EPI) benefits from the proposed shear layer as well as the downscaling structure. (a) Ground truth and input EPIs; (b) EPI without shear layer; (c) EPI without the learning-based Laplacian pyramid structure; (d) EPI with our proposed pipeline.

to a 2D EPI but also to a feature map in the network, and thus, is embedded into the neural network for end-to-end optimization. In contrast, the shear operation in [31] is only applied for pre- and post-processing of the 2D EPI; 2) We predict multiple high angular resolution EPIs, one for each shear amount, and blend them in the high-level feature domain (please refer to Sec. 4.2). While the method in [31] feeds a single sheared EPI into the network, and blends the reconstructed EPIs through weighted average.

In the following, we introduce the reconstruction net that takes the sheared EPI $f_h(E_{LR}; \alpha_h)$ as input, where a learning-based Laplacian pyramid is proposed to implement the downscaling operation described in Sec. 3.2. Specifically, we first adopt three convolutional operations with strides [4, 2, 1] in the spatial dimension (i.e., α_u = [4, 2, 1]) to construct a multi-scale structure; then upsample the first and the second feature layers using deconvolution with stride 2; finally, we construct residual layers by subtracting each upsampled feature from the corresponding layer in the multi-scale structure. Please refer to Fig. 3 (top right) for the network connection. The proposed learningbased Laplacian pyramid owns the following superiority: 1) a downscaled feature layer is provided for addressing the aliasing issue echoing with the analysis in Sec. 3.2; 2) a learning-based representation of low-frequency and high-frequency for extracting the aliasing information. As illustrated in Fig. 5 (c) and (d), without the proposed pyramid structure, the reconstructed EPI (5 (c)) shows obvious aliasing artifacts.

For the **prefilter layer**, we use 1D convolutional operations with different kernel sizes for different spatial scales. We use the Gaussian function to initialize the trainable weights. Moreover, we adopt multiple shape parameters in each layer to achieve an effect of different degrees of antialiasing, which enables the network for addressing different degrees of aliasing. Specifically, we extend the Gaussian kernel described in Eqn. 4 to a higher dimension as follows

$$\theta_p(u,c) = \frac{\exp(-\frac{u^2}{2\sigma^2(c)})}{\sum_u \exp(-\frac{u^2}{2\sigma^2(c)})},$$

where θ_p is the 1D convolutional kernel in the prefilter layer and $\sigma(c)$ is the shape parameter for channel *c*. For each spatial scale in the prefilter layer, the maximum value of the shape parameter σ_{max} is selected according to the theoretical analysis in Sec. 3, and $\sigma(c)$ for each channel *c* is uniformly distributed from zero to σ_{max} . To be consistent with the definition of filtering, bias and activation function are not applied in the prefilter layer¹.

Straight after the prefilter layer, we apply a convolutional layer for both the spatial and the angular dimensions, and a deconvolution layer (spatial) for upscaling the feature maps to the original spatial resolution. Then the three streams are concatenated together in the channel dimension, being followed by a convolutional layer with batch normalization [32]. The last layer is a deconvolution layer (angular) that is designed to reconstruct high angular resolution EPI (feature maps) through upsampling. Each layer in the reconstruction net except for the prefilter layer and the angular deconvolution layer is followed by a leaky ReLU. Please refer to the supplementary for detailed specification of the reconstruction net.

In summary, the entire reconstruction module f_R can be described as follows

$$f_R(E_{LR};\theta_r) = \{ f_h(f_r(f_h(E_{LR};\alpha_h^{(i)});\theta_r); -\frac{\alpha_h^{(i)}}{\alpha_s}) \}, \quad (7)$$

(i)

where $\alpha_h^{(i)}$ is the *i*-th shear amount and $f_r(\cdot; \theta_r)$ denotes the reconstruction net with trainable parameters θ_r . The shear amount α_h and angular upsampling factor α_s are two hyperparameters, set as $\alpha_h \in [-9, -6, -3, 0, 3, 6, 9]$ and $\alpha_s = 3, 4$. Note that the shear values can be readjusted without retraining, and the reconstruction scale is flexible via cascade.

4.2 Fusion module

The fusion module serves to blend multiple feature layers with different shear amounts, namely different candidates of the reconstructed EPI, as shown in Fig. 3. Note that only the candidate sheared by the optimal or close to the optimal rendering disparity will get a high quality reconstruction result. In this way, the overall framework is able to select optimal rendering disparity without explicit disparity estimation.

The input of the fusion module is a 3D tensor, i.e., a high angular resolution feature map, rather than a 2D EPI slice, to prevent unnecessary feature extractions from 2D EPIs. We first concatenate the sheared layers yielded by the reconstruction module along the channel dimension. The last step contains a fusion net that produces the final high angular resolution EPI by blending the multiple sheared feature layers together. The proposed fusion net follows the U-net structure, while it varies in terms of: 1) a convolutional layer with kernel size [1,1] is set in the front of the fusion net to compress the feature maps; 2) the encoder (decoder) downscales (upscales) the feature maps only in the spatial dimension. Each layer except for the last layer is followed by a leaky ReLU. Please refer to the supplementary for the detailed architecture of the fusion net.

Combining with the reconstruction module described in Eqn. 7, the proposed DA²N with all the trainable parameters can be represented as

$$f_D(E_{LR};\boldsymbol{\theta}) = f_F(f_R(E_{LR};\boldsymbol{\theta}_r);\boldsymbol{\theta}_f), \qquad (8)$$

where $\boldsymbol{\theta} = \{\theta_r, \theta_f\}$, and $f_F(\cdot; \theta_f)$ is the fusion net with trainable parameters θ_f .

1. After the end-to-end training, the filters in the prefilter layers are no longer simple Gaussian functions but are still some kinds of lowpass filters. Please refer to the supplementary for more details.



(b) EPIs from regular light fields and their Fourier spectrums

Fig. 6. Pseudo EPIs and regular EPIs in non-Lambertian scenes can be described by curves and variation of pixel intensity (color). These two EPIs show the same property in the Fourier domain as well.

5 NETWORK TRAINING

5.1 Training dataset

A sufficiently large set of LFs with plenty views for extracting large-disparity LFs or with non-Lambertian property is favored for the training procedure. Unfortunately, commonly-used LF dataset (LFs from HCI datasets [16] or a Lytro Illum [33]) is lack of data with non-Lambertian effect. Meanwhile, existing LF datasets suffer small disparity range and limited angular resolution. For example, in the HCI datasets, the disparity range is around ± 2 pixel and angular resolution is 9×9 . For the Lytro LF, the disparity is nearly within 1 pixel range, and the valid views are 8×8 . With an angular downsampling factor of 3, these datasets only provide an input with around 6 and 3 disparity ranges, respectively.

Considering the above problems, we adopt EPIs from regular (calibrated) LFs for training from scratch and **pseudo EPIs from unstructured LFs** for fine-tuning. Our key insight is that pseudo EPIs exactly accord with the non-Lambertian property in regular EPIs. In the following, we explain this insight by expanding the Fourier analysis in Sec. 3.1 to the pseudo EPIs from unstructured LFs.

An unstructured LF may be generated from an irregularly moving camera, including a non-linear movement along the main angular dimension², a slight shaking along the other angular dimension or around the normal of the camera plane. The former one introduces "curves" in the recorded EPI, as shown in the top of Fig. 6(a). While the latter one leads to a sporadic "appearance" of certain pixels along the angular dimension, i.e., abrupt variation of intensity (or color), as shown in the bottom of Fig. 6(a). In the first case, the curve in a pseudo EPI can be split into lines with a series of different disparities. Therefore, its spectrum will be $D\Omega_u + \Omega_s = 0$, where D is the set of disparities from the curve (as shown in the top right spectrum of Fig. 6(a)). In the second case, the intensity (color) variation can be approximated by a band-limited signal [6], thus, its spectrum is $d\Omega_u + \Omega_s = \pm \beta'/Z$, where β' is the band width of the signal. By combining these two case, we then have the Fourier spectrum described as $D\Omega_u + \Omega_s = \pm \beta'/Z$ (see the bottom right spectrum of Fig. 6(a)). In a non-Lambertian scene, the regular EPI can be also regarded as the combination of curves and variation of pixel intensity,

which can be clearly seen in Fig. 6(b). And the spectrum of the EPI also conforms to the formula $D\Omega_u + \Omega_s = \pm \beta'/Z$ (as shown in the right spectra of Fig. 6(b)).

In addition, unstructured LFs can easily provide a high angular resolution (up to 600 pixels in our dataset) due to the acquisition method (e.g., by a hand-hold video camera), therefore, is able to offer EPIs with a large disparity. In conclusion, the pseudo EPI dataset provides the same non-Lambertian description and a better data abundance for the network training.

5.2 Objective

Based on the high angular resolution EPI reconstruction pipeline defined in Eqn. 8, the objective is to minimize a loss function $\mathcal{L}(\cdot)$ between the high angular resolution EPI label E_{HR} and the reconstructed EPI $f_D(E_{LR}; \theta)$ as

$$\min_{\boldsymbol{\theta}} \sum_{\langle E_{LR}, E_{HR} \rangle} \mathcal{L}(f_D(E_{LR}; \boldsymbol{\theta}), E_{HR}),$$

where $\langle E_{LR}, E_{HR} \rangle$ is the training set of low and high angular resolution EPI tuples.

We combine two terms in the loss function \mathcal{L} , a L_1 distance and a perceptual loss (also known as high-level feature matching loss) between the label EPI and the reconstructed EPI. The perceptual loss can be computed from part of the feature layers in the autologous network [35] or other pre-trained networks [36], e.g., the commonly-used VGG-16 [37]. To prevent loss fluctuation in the training procedure, we adopt a pre-trained network is an encoder-decoder with a similar structure of the fusion net. While the main differences are that no skip connections are used and the downscales (upscales) are performed in both the spatial and angular dimensions. We apply the same training dataset as introduced in Sec. 5.1 for the encoder-decoder. The loss function is then defined as

$$\mathcal{L}(\bar{E}_{HR}, E_{HR}) = \|\bar{E}_{HR} - E_{HR}\|_{1} + \sum_{l} \lambda_{feat}^{(l)} \|\phi^{(l)}(\hat{E}_{HR}) - \phi^{(l)}(E_{HR})\|_{1},$$
⁽⁹⁾

where $\{\phi^{(l)}\}\$ is a set of layers in the encoder including conv1_1, conv2_1 and conv3_1, and $\lambda_{feat} = [0.2, 0.2, 0.1]$ is a set of hyperparameters for the perceptual loss.

5.3 Implementation details

The upsampling scale α_s in the angular dimension is fixed for the training, but the reconstruction scale is flexible by combining model cascade and bicubic downsampling (angular). For instance, for interpolation ×7, we will first achieve a ×9 reconstruction with two cascades ($\alpha_s = 3$) and then downsample the resulting LF to the desired angular resolution. The training is performed on the Y channel (i.e., the luminance channel) of the YCbCr color space. To speed up the training procedure, we apply patch-wise training strategy by sampling sub-EPI pairs in the training data, and then applying mini-batches of size 28. We use zero padding in each layers according to the filter size.

Instead of training the DA²N using mixed datasets from the regular LFs and the unstructured LFs, we divide the training into the following two steps.

^{2.} In this paper, we only consider a 3D unstructured LF, namely, the LF is recorded along a main angular dimension.



Our reconstructed view (PSNR / SSIM)

Ours (37.99/0.9915)

Ground truth

Fig. 7. Comparison of the results ($\times 16$ upsampling) on the LFs from the ICME DSLF dataset [34]. The results show one of our reconstructed view, EPIs extracted from LFs reconstructed by each methods. The PSNR and SSIM values are averaged on the reconstructed views (input views are excluded).

| | TABLE 2 | |
|----------------------------------|--------------------------------------|----------------------------|
| Quantitative results (PSNR/SSIM) | of reconstructed LFs on the LFs from | the ICME DSLF dataset [34] |

| Method | | DD1 | | DD2 | | DD3 | | Average |
|----------|-----------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | | (7px) ×8 | (14px) ×16 | (7px) ×8 | (14px) ×16 | (7px) ×8 | (14px) ×16 | 0 |
| | DIBR [39], [40] | 39.93/0.9884 | 39.31/0.9876 | 37.99/0.9743 | 37.88/0.9707 | 36.09/0.9860 | 35.15/0.9786 | 37.73/0.9809 |
| | Kalantari et al. [10] | 45.66/0.9901 | 42.97/0.9840 | 38.95/0.9712 | 35.94/0.9631 | 35.88/0.9741 | 31.92/0.9593 | 38.55/0.9736 |
| Baseline | Wu et al. [12] | 44.81/0.9905 | 42.31/0.9761 | 40.58/0.9833 | 37.90/0.9739 | 37.83/0.9931 | 34.87/0.9794 | 39.72/0.9827 |
| | Yeung et al. [13] | 47.10/0.9965 | 41.36/0.9836 | 39.61/0.9915 | 35.45/0.9834 | 39.96/0.9951 | 32.83/0.9749 | 39.39/0.9875 |
| | Wu et al. [31] | 46.89/0.9899 | 44.81/0.9860 | 40.87/0.9716 | 38.65/0.9668 | 37.96/0.9859 | 35.83/0.9771 | 40.84/0.9796 |
| | w/o shear | 46.78/0.9948 | 44.05/0.9930 | 41.55/0.9922 | 39.32/0.9903 | 40.08/0.9951 | 36.68/0.9903 | 41.41/0.9926 |
| | w/o downscale | 46.71/0.9964 | 44.49/0.9948 | 41.38/0.9921 | 39.22/0.9893 | 39.79/0.9942 | 36.41/0.9890 | 41.33/0.9926 |
| Ablation | w/o prefilter | 47.92/0.9972 | 45.87/0.9949 | 42.42/0.9923 | 40.52/0.9912 | 41.08/0.9953 | 37.51/0.9897 | 42.55/0.9934 |
| | w/o pseudo EPIs | 46.53/0.9962 | 41.21/0.9834 | 41.93/0.9920 | 38.14/0.9862 | 40.58/0.9954 | 33.60/0.9781 | 40.33/0.9886 |
| | Our proposed | 48.23/0.9971 | 46.19/0.9956 | 42.65/0.9934 | 40.77/0.9915 | 41.38/0.9960 | 37.99/0.9915 | 42.87/0.9942 |

Ablation studies are performed for the network without the shear layers ("w/o shear" for short), without the Laplacian pyramid structure for downscaling ("w/o downscale" for short), without prefilter layers ("w/o prefilter" for short) and without pseudo EPIs from unstructured LFs for fine-tuning ("w/o pseudo EPIs" for short). The disparity range ($d_{max} - d_{min}$) between the neighboring views in the input light field is listed in the table.

Training from scratch on regular LFs. In this step, about 13K EPIs are extracted from 11 LFs from the Stanford Light Field Archive [41] with stride 20 in the spatial dimensions. In our implementation, we extract sub-EPIs of size 6×72 for the input examples and 16×72 for the labels from the EPIs with stride 16 in the spatial dimensions, and donwnsample the angular dimension using nearest sampling. Through this way, we obtain about 725K examples from the structured LFs. We initialize the weights of both convolution and deconvolution layers by drawing randomly from a Gaussian distribution with a zero mean and standard deviation 1×10^{-3} , and the biases by zero. The prefilter layers are initialized with Gaussion function as mentioned in Sec. 4.1. We use ADAM solver [42] as the optimization method with learning rate of 1×10^{-5} for the prefilter layers and 4×10^{-5} for the rest of the layers, and $\beta_1 = 0.9$, $\beta_2 = 0.999$. The network converges after 600K steps of backpropagation.

Fine-tuning on unstructured LFs. In this step, we extract 1888 pseudo EPIs (from 11 unstructured LFs by Yücer *et al.* [43]) at 600×1280 (or 1920) (angular×spatial) resolution with stride 31/400 along dimension v/t. About 1400K sub-EPIs with resolution 31×72 (for the labels) are extracted

from pseudo EPIs (stride 92/23 for the angular/spatial dimension). The learning rate is set to 1×10^{-6} for the prefilter layers and 1×10^{-5} for the rest of the layers. The network converges after 800K steps of backpropagation. The training model is implemented using the *Tensorflow* framework [44]. The entire training steps (training from scratch and finetuning) takes about 50 hours on a NVIDIA TITAN Xp.

6 EVALUATION AND APPLICATION

In this section, we evaluate the proposed DA²N on several datasets, including LFs from gantry system, LFs from a hand-held plenoptic camera (Lytro Illum [33]) and LFs from camera array system. Note that all these datasets are not overlap with our training data in terms of acquisition geometry. The term "large disparity" is defined as the maximum disparity between neighboring views being larger than 5 pixels or the disparity range being more than 10 pixels.

We mainly compare our approach with four state-ofthe-art methods, including two depth and learning-based methods by Kalantari *et al.* [10] and Wu *et al.* [31], and two depth independent and learning-based methods by Yeung *et al.* [13] and Wu *et al.* [12]. In addition, we perform a



Fig. 8. Comparison of the results ($\times 16$ upsampling) on the LFs from the MPI Light Field Archive [38]. The green lines in the ground truth EPI are input views. The PSNR and SSIM values are averaged on the reconstructed views (input views are excluded).

TABLE 3 Quantitative results (PSNR/SSIM) of reconstructed LFs ($\times 16$) on the LFs from the MPI Light Field Archive [38].

| Bikes (23.5px) | FairyCollection (33.5px) | LivingRoom (14px) | Mannequin (14.5px) | WorkShop (23px) | Average |
|----------------|---|---|---|---|--|
| 30.03/0.9525 | 30.16/0.9276 | 44.07/0.9860 | 36.07/0.9613 | 32.33/0.9690 | 34.53/0.9593 |
| 27.89/0.9101 | 28.37/0.9231 | 39.95/0.9901 | 36.12/0.9689 | 35.72/0.9677 | 33.61/0.9520 |
| 27.23/0.9045 | 27.69/0.9069 | 40.87/0.9909 | 35.66/0.9748 | 30.08/0.9371 | 32.31/0.9428 |
| 30.16/0.9365 | 30.37/0.9427 | 43.10/0.9896 | 36.74/0.9792 | 33.81/0.9561 | 34.84/0.9608 |
| 31.58/0.9425 | 35.02/0.9721 | 44.56/0.9885 | 38.21/0.9760 | 36.47/0.9796 | 37.17/0.9717 |
| 35.79/0.9844 | 36.23/0.9813 | 45.91/0.9963 | 40.83/0.9918 | 40.11/0.9936 | 39.77/0.9895 |
| | Bikes (23.5px) 30.03/0.9525 27.89/0.9101 27.23/0.9045 30.16/0.9365 31.58/0.9425 35.79/0.9844 | Bikes (23.5px) FairyCollection (33.5px) 30.03/0.9525 30.16/0.9276 27.89/0.9101 28.37/0.9231 27.23/0.9045 27.69/0.9069 30.16/0.9365 30.37/0.9427 31.58/0.9425 35.02/0.9721 35.79/0.9844 36.23/0.9813 | Bikes (23.5px)FairyCollection (33.5px)LivingRoom (14px)30.03/0.952530.16/0.927644.07/0.986027.89/0.910128.37/0.923139.95/0.990127.23/0.904527.69/0.906940.87/0.990930.16/0.936530.37/0.942743.10/0.989631.58/0.942535.02/0.972144.56/0.9885 35.79/0.984436.23/0.981345.91/0.9963 | Bikes (23.5px)FairyCollection (33.5px)LivingRoom (14px)Mannequin (14.5px)30.03/0.952530.16/0.927644.07/0.986036.07/0.961327.89/0.910128.37/0.923139.95/0.990136.12/0.968927.23/0.904527.69/0.906940.87/0.990935.66/0.974830.16/0.936530.37/0.942743.10/0.989636.74/0.979231.58/0.942535.02/0.972144.56/0.988538.21/0.9760 35.79/0.984436.23/0.981345.91/0.996340.83/0.9918 | Bikes (23.5px)FairyCollection (33.5px)LivingRoom (14px)Mannequin (14.5px)WorkShop (23px)30.03/0.952530.16/0.927644.07/0.986036.07/0.961332.33/0.969027.89/0.910128.37/0.923139.95/0.990136.12/0.968935.72/0.967727.23/0.904527.69/0.906940.87/0.990935.66/0.974830.08/0.937130.16/0.936530.37/0.942743.10/0.989636.74/0.979233.81/0.956131.58/0.942535.02/0.972144.56/0.988538.21/0.976036.47/0.9796 35.79/0.984436.23/0.981345.91/0.996340.83/0.991840.11/0.9936 |

The disparity range $(d_{\max} - d_{\min})$ between the neighboring views in the input light field is listed in the table.

series of ablation studies of our approach. The quantitative evaluations are reported by measuring the average PSNR and SSIM [45] values over the synthesized views (in both figures and tables). Please also refer to the submitted supplementary video for more qualitative results.

6.1 Light fields from gantry system

We use the dataset from the ICME 2018 Grand Challenge on Densely Sampled Light Field Reconstruction [34] ("ICME D-SLF dataset" for short) and the MPI Light Field Archive [38], for the evaluation of LFs from gantry system. In this experiment, we reconstruct a 1×193 LF using 13 views for the ICME DSLF dataset [34] and a 1×97 LF using 7 views for the MPI Light Field Archive [38], i.e., $\times 16$ upsampling.

In this experiment, the baseline methods are a conventional depth image-based rendering (DIBR) approach (the approach by Jeon *et al.* [39] for depth estimation and the approach by Chaurasia *et al.* [40] for view synthesis), the depth independent and learning-based approach by Wu *et al.* [12] and Yeung *et al.* [13], and the depth and learningbased approach by Kalentary *et al.* [10] and Wu *et al.* [31]. The baseline methods [10], [13] are retrained for better comparison.

Fig. 7 shows the results (\times 16 upsampling) of the case *DD2* and *DD3* from the ICME DSLF dataset [34], with disparity ranges between neighboring views [-2, 12] and

[-8, 6], respectively. Both cases contain challenging reflective surfaces, e.g., the ladder in the *DD2* and the metal bottle in the *DD3*. The depth-based view synthesis approach (DIBR [39], [40]) shows tearing artifacts around the occluded thin structures, as shown in the EPIs of the case *DD2*. It also fails to produce reasonable result around the reflective surfaces due to the ambiguous depth in those regions, as shown in the EPIs of the case *DD3*. The results by Wu *et al.* [12] shows ghosting or blurring effects when addressing the severe aliasing effects in these two cases. Among the baseline methods, the proposed DA²N produces plausible results when reconstructing the LFs with both large disparity and non-Lambertian effect.

Fig. 8 shows the results (×16 upsampling) of the case *FairyCollection* and *Mannequin* from the MPI Light Field Archive [38], with disparity ranges between neighboring views [-16.5, 17] and [-7.5, 7], respectively. In the demonstrated non-Lambertian materials (EPIs), the disparities are about 14 and 6 pixels, respectively. The significant disparity as well as the non-Lambertian objects (refractive glass) in *Fairycollection* brings extensive challenge for LF reconstruction. Similarly, the specular board in *Mannequin* also suffers non-Lambertian effect. As clearly demonstrated by both the quantitative and qualitative results, the proposed DA²N succeeds to reconstruct the distorted curves by the refraction of the glass. Meanwhile, it can effectively reconstruct not only the white specularity but also the covered red hand-



Fig. 9. Comparison of the results on the LFs from Lytro Illum (\times 3 upsamping). The results show one of our reconstructed views, close-up versions of the image portions in the yellow and blue boxes, and the EPIs located at the red line shown in the left images. We also show a close-up of the portion of the EPIs in the red box. LFs are from the *30 Scenes* set [10] and the *Refractive and Reflective Surfaces* category [46].

| TABLE 4 |
|---|
| Quantitative results (PSNR/SSIM) of reconstructed views on the LFs from plenoptic cameras |

| Method | 30 Se | 30 Scenes | | Reflective & Refractive Surfaces | | Occlusions | |
|------------------------------|--------------|----------------------|--------------|----------------------------------|--------------|--------------|--------------|
| | $\times 3$ | $\times 7$ | $\times 3$ | $\times 7$ | $\times 3$ | $\times 7$ | 0 |
| Kalentary <i>et al.</i> [10] | 39.62/0.9782 | 38.21/0.9736 | 37.78/0.9713 | 35.84/0.9416 | 34.02/0.9554 | 31.81/0.8945 | 35.61/0.9474 |
| Wu et al. [12] | 41.02/0.9875 | 36.28/0.9648 | 41.71/0.9886 | 36.48/0.9619 | 38.30/0.9701 | 32.36/0.9240 | 37.28/0.9628 |
| Yeung et al. [13] | 44.53/0.9900 | 39.22 /0.9773 | 42.56/0.9745 | 36.47/0.9472 | 39.27/0.9450 | 32.68/0.9061 | 38.54/0.9510 |
| Wu et al. [31] | 42.46/0.9861 | 38.21/0.9711 | 42.28/0.9773 | 36.55/0.9534 | 38.11/0.9443 | 32.19/0.9067 | 37.73/0.9509 |
| Our proposed | 43.69/0.9949 | 38.99/ 0.9863 | 43.25/0.9906 | 36.72/0.9751 | 39.82/0.9706 | 33.14/0.9504 | 38.76/0.9748 |

The metrics are averaged over 113 LFs.

writing. The quantitative results on real-world LFs in the MPI Light Field Archive [38] is listed in Table 3.

Ablation study. We validate the overall strategies, including the shearing, downscaling and prefiltering operations, and the adopted pseudo EPIs for fine-tuning, by performing the following ablation studies: without the shear layers ("w/o shear" for short), without the Laplacian pyramid structure for downscaling ("w/o downscale" for short), without prefilter layers ("w/o prefilter" for short) and without pseudo EPIs from unstructured LFs for finetuning ("w/o pseudo EPIs" for short). The quantitative results is listed in Table 2. Due to the incompleteness of the proposed shearing, downscaling and prefiltering operations for aliasing handling, these approaches show performance degradation in various degrees. Moreover, as indicated by Table 2, the performance gap between the network trained on regular LFs ("w/o pseudo EPIs") and the fine-tuned network ("our proposed") turns bigger for larger downsampling rate, i.e., from 1.07dB for $\times 8$ downsampling rate to 3.96dB for $\times 16$ downsampling rate. In other words, the network without fine-tuning on the pseudo EPIs from unstructured LFs suffers a significant performance decrease, especially when addressing severe aliasing.

6.2 Light fields from Lytro Illum

We evaluate the proposed approach using the *30 Scenes* dataset by Kalantari *et al.* [10] (contains 30 LFs) , and the *Refractive and Reflective Surfaces* category (contains 32 LFs) as well as the *Occlusions* category (contains 51 LFs) from the Stanford Lytro Light Field Archive [46]. In this experiment, we reconstruct a 7×7 LF from 3×3 views (×3 upsampling)

and a 8×8 LF from 2×2 views (×7 upsampling). The baseline methods are two depth independent approaches by Wu *et al.* [12] and Yeung *et al.* [13], and two depth-based approaches by Kalantari *et al.* [10] and Wu *et al.* [31]. Since the baseline methods by Wu *et al.* [12], [31] used the same structured LFs as the proposed DA²N for training and the other two methods are trained specifically on LFs from Lytro Illum, they are not retrained or fine-tuned in this evaluation.

Fig. 9 shows the results of the case IMG1743 from the 30 Scenes set [10] and Reflective-13 from the Refractive and *Reflective Surfaces* category [46] using a 3×3 input for reconstructing a 7×7 LF. The two cases contain large disparities for better demonstrating the aliasing problem. In both cases, the approach by Kalantari et al. [10] produces tearing artifacts near the occlussion boundaries of the foreground objects, which can be clearly shown in the close-up images and EPIs. The approach by Yeung et al. [13] fail to address the aliasing effects around the backgrounds, where the maximum disparity is about 13 pixels. In comparison, our approach provides better results in terms of both local structure (shown in the close-up images) and overall coherency (shown in the EPIs). Table 4 lists the quantitative results on the evaluated LFs from Lytro Illum. The proposed DA²N achieves the highest average PSNR and SSIM values compared with other approaches. Note that the DA²N is not trained on LFs from Lytro Illum and we do not require additional retraining or fine-tuning for the $\times 3$ and $\times 7$ upsampling scales.

Model size and running time. The model size with respect to number of trainable parameters and running time (with GPU acceleration) for the reconstruction of LFs from Lytro Illum are compared in Table 5. It can be shown that the



Fig. 10. Comparison of the results on the LFs from camera array system (×3 upsamping). The results show one of the synthesized result, a reference image from the input LF, a portion of the EPI extracted from the location marked in the red line.

TABLE 5 Comparison of model size (number of trainable parameters) and running time for the reconstruction of LFs from Lytro Illum.

| Method | Model size | Running time (s) |
|--------------------------|------------|------------------|
| Kalantari et al. [10] | 1654004 | 532.5 |
| Wu et al. [12] | 56769 | 303.2 |
| Yeung <i>et al.</i> [13] | 392828 | 6.3 |
| Wu et al. [31] | 206305 | 88.3 |
| Our proposed | 541827 | 19.8 |

efficiency of our method is comparable as Yeung *et al.* [13], yet the performance regarding to the aliasing issue of our method is superior as demonstrated in Fig. 9.

6.3 Light fields from camera array system

In this section, the camera array-based dataset provided by Wang *et al.* [47] and the microscope dataset provided by Lin *et al.* [48] are evaluated. We reconstruct a LF with 7×7 views from a 3×3 input (×3 upsampling).

Fig. 10 shows the reconstructed LFs of the case Scene 4 from the dataset by Wang et al. [47] and the case worm from the microscope dataset by Lin et al. [48]. The first case [47] shows a simple scene with clear occlussion relations and maximum disparity around 15 pixels. The results produced by the baseline methods from Kalantari et al. [10] and Yeung et al. [13] show aliasing effect due to the large disparity, which can be clearly seen in the sub-aperture images and the EPIs in Fig. 10. While the approach by Wu et al. [12] shows blurring artifects due to the high-frequency loss introduced by their "blur-restoration-deblur" framework. The second case shows a translucent worm (drosophila larva) captured by a camera array-based microscopy [48], which is challenge due to the large disparity as well as the non-Lambertian. Due to the ambiguous depth information, the approach by Kalantari et al. [10] fails to render views with reasonable appearance. The result by Wu et al. [12] shows blur effects due to the large disparities. While the approach by Yeung *et* al. [13] produces views with aliasing effects, as shown by the EPI marked with the red arrow in Fig. 10. Among these approaches, our DA²N is able to reconstruct a high quality LF in this non-Lambertian scene.

6.4 Application for view extrapolation

We further extend the proposed LF reconstruction approach for view extrapolation. Different with the LF interpolation, however, the view extrapolation requires the prediction of occluded regions which are hidden from all the input views. In the EPI domain, it is actually closer to the outpainting problem, that is, the network needs to search texture patterns from the input (or reconstructed) EPI and pads them outwards. The proposed DA²N ("shearing-downscaling-prefiltering" structure) is designed to solve the aliasing problem by reducing disparity (range) in the sampled light field, making it easier for the network to find the correct patterns for view extrapolation. In addition, by combining with a discrimination network [49], the DA²N is able to inference the extrapolated view and predict the occluded background simultaneously in the EPI domain.

Implementation. The overall network architecture for light field view extrapolation is the same as the one we introduced in Sec. 4. The main difference is that we introduce a discrimination network to construct a generative adversarial network (GAN) together with the proposed DA²N. The discrimination network has a similar structure with that introduced by Iizuka et al. [50], which contains a global discriminator and a local discriminator for both global and local consistency in the EPI domain. Specifically, the output of the generator (DA^2N) is fed to the global discriminator, while the synthesized part is fed to the local discriminator. Then a concatenation layer is applied to connect the outputs from the two discriminators in the channel dimension and yields the final discrimination value. Accordingly, we also add a standard adversarial term [49] to the original loss function described in Eqn. 9 for the generator

$$\mathcal{L}_{gen}(\vec{E}, E) = \mathcal{L}(\vec{E}, E) + \lambda_{adv} \log(f_{dis}(\vec{E})),$$

where E and E are the extrapolated EPI and its corresponding label, $f_{dis}(\cdot)$ is the discrimination network and $\lambda_{adv} = -5 \times 10^{-4}$ is a weighing hyperparameter.

In this application, the network is trained on the regular LFs from the Stanford Light Field Archieve [41]. Sub-EPIs with size 14×64 are extracted as the training labels, where 7 views are used to extrapolate the rest 7 views, i.e., $\times 2$ extrapolation. Note that the extrapolation scale is also flexible similar with the view interpolation via cascade. About 1424K examples are extracted from the 11 LFs [41]. The learning rate is set to 3×10^{-4} for the generator and 6×10^{-5} for the discrimination network. The network converges after 800K steps of backpropagation.



Fig. 11. Comparison of the extrapolation results on the LFs from gantry system. In each case, we applies 2 images (marked in the red lines in the ground truth EPIs) to render a 1×35 LF with 7 interpolated views and 28 extrapolated views. The PSNR and SSIM metrics are averaged on the extrapolated views.

Results. We compared the proposed view extrapolation approaches with two baseline methods, a DIBR [39], [40] approach combined with learning-based image inpainting [50] for occluded regions and a multiplane image-based view extrapolation approach by Zhou et al. [24]. Fig. 11 shows the extrapolation results for the LF from the gantry system [34]. In each case, a LF with 2 horizontal views (marked by the red lines in the ground truth EPIs in Fig. 11) are applied to synthesize a 1×35 LF, where 7 views are from the interpolation and 28 views from the extrapolation, i.e., $\times 5$ extrapolation. The extrapolated views produced by the inpaintingbased approach [50] appear severe distortion around the occluded regions due to the unreliable estimation of the disoccluded background, which can be shown by the EPIs in Fig. 11. While the results produced by the multiplane imagebased approach [24] show ghosting artifacts due to the inaccurate depth estimation around occlusion boundaries, which is then magnified by the extrapolation. In comparison, the proposed DA²N is able to extrapolate high quality LF in terms of both visual coherency and quantitative measurement (see the PSNR and SSIM values averaged on the extrapolated views in Fig. 11).

6.5 Limitations

A main limitation of the proposed anti-aliasing framework is that we could not reconstruct the non-Lambertian effect that is sampled under an extremely sparse sampling. As demonstrated in Fig. 12(a), when the downsampling rate in the angular dimension is greater than $\frac{\pi Z}{\beta}$ + 1, the expansionary spectra will overlap with its replicas (please refer to Eqn. 1 and Fig. 2). Unfortunately, this overlap or aliasing part cannot be eliminated through downscaling or prefiltering in the spatial dimension. Another limitation is that the proposed DA²N sometimes fail to reconstruct regions with repetitive pattern. As illustrated in Fig. 12(b), the reconstruction net generates multiple "plausible" results using different shear amounts, which lead to ambiguity after being fed into the subsequent fusion net.



Fig. 12. Limitations of the proposed method. (a) In a non-Lambertian case with extremely sparse input views (top left), there will be overlap between the expansionary spectra and its replicas (right), which cannot be eliminated through downscaling or prefiltering in the spatial dimension; (b) In some cases with repetitive pattern, the reconstruction net will generate several "plausible" outputs using different shear amounts (right), which will lead to ambiguity when being blended together.

7 CONCLUSION AND FUTURE WORK

Solving the aliasing problem is a key issue for LF reconstruction under the large disparity and non-Lambertian challenges. In this paper, we introduce a deep learning pipeline that is specifically designed for aliasing handling. We theoretically show that the internal components of the proposed deep learning pipeline is equivalent to or even superior to the conventional reconstruction filter in the Fourier domain in solving the aliasing problem. Specifically, the network includes a shear layer to warp the input LF to certain depths, a learning-based Laplacian pyramid to reduce disparity range and a prefilter layer to further remove the aliasing component. To explicitly enhance the capacity on aliasing handling, we begin by training the network from scratch using regular LFs and then fine-tune it with pseudo EPIs from unstructured LFs. We analytically indicate that pseudo EPIs have the same non-Lambertian property in terms of the visual features and the aliasing issue. By performing evaluations on LFs from different acquisition geometries (gantry system, hand-held plenoptic camera and

camera array system), we have demonstrated the efficacy and robustness for aliasing handling.

An interesting aspect of future work is to promote the fusion net with flexible input, that is, a variable number of shear layers. This setting will enable users to customize the network architecture according to the disparity range as well as the computational complexity without retraining. To achieve this feature, the fusion net is required to be fully convolutional along the angular, width and shear dimensions, i.e., a 3D network, while keeping only the angular and width dimensions before producing the final EPI.

ACKNOWLEDGMENTS

This work was supported by the Major Program of National Natural Science Foundation of China NSFC No.61991400 and No.61991401, Science and Technology Major Projects of Liaoning Province No.2020JH1/10100008, NSFC No.61827805, No.61531014, No.61861166002 and No.6181001011, and Fundamental Research Funds for the Central Universities No. 100802004.

REFERENCES

- M. Levoy and P. Hanrahan, "Light field rendering," in *Proceedings* of the 23rd annual conference on Computer graphics and interactive techniques. ACM, 1996, pp. 31–42.
- [2] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen, "The lumigraph," in Siggraph, vol. 96, no. 30, 1996, pp. 43–54.
- [3] B. Wilburn, N. Joshi, V. Vaish, E.-V. Talvala, E. Antunez, A. Barth, A. Adams, M. Horowitz, and M. Levoy, "High performance imaging using large camera arrays," in ACM TOG, vol. 24, no. 3. ACM, 2005, pp. 765–776.
- [4] R. S. Överbeck, D. Erickson, D. Evangelakos, M. Pharr, and P. Debevec, "A system for acquiring, processing, and rendering panoramic light field stills for virtual reality," in *SIGGRAPH*, 2019.
 [5] J.-X. Chai, X. Tong, S.-C. Chan, and H.-Y. Shum, "Plenoptic sam-
- [5] J.-X. Chai, X. Tong, S.-C. Chan, and H.-Y. Shum, "Plenoptic sampling," in *Proceedings of the 27th annual conference on Computer* graphics and interactive techniques. ACM Press/Addison-Wesley Publishing Co., 2000, pp. 307–318.
- [6] C. Zhang and T. Chen, "Spectral analysis for sampling imagebased rendering data," *IEEE Transactions on Circuits and Systems* for Video Technology, vol. 13, no. 11, pp. 1038–1050, 2003.
- [7] S. Vagharshakyan, R. Bregovic, and A. Gotchev, "Light field reconstruction using shearlet transform," *IEEE TPAMI*, vol. 40, no. 1, pp. 133–147, 2018.
- [8] J. Stewart, J. Yu, S. Gortler, and L. McMillan, "A new reconstruction filter for undersampled light fields," in ACM International Conference Proceeding Series. Eurographics Association/Association for Computing Machinery, 2003.
- [9] J. Flynn, I. Neulander, J. Philbin, and N. Snavely, "Deepstereo: Learning to predict new views from the world's imagery," in *CVPR*, 2015.
- [10] N. K. Kalantari, T.-C. Wang, and R. Ramamoorthi, "Learningbased view synthesis for light field cameras," ACM TOG, vol. 35, no. 6, 2016.
- [11] Y. Yoon, H.-G. Jeon, D. Yoo, J.-Y. Lee, and I. So Kweon, "Learning a deep convolutional network for light-field image superresolution," in CVPRW, 2015, pp. 24–32.
- [12] G. Wu, Y. Liu, L. Fang, Q. Dai, and T. Chai, "Light field reconstruction using convolutional network on epi and extended applications," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 7, pp. 1681–1694, 2018.
- [13] H. W. F. Yeung, J. Hou, J. Chen, Y. Y. Chung, and X. Chen, "Fast light field reconstruction with deep coarse-to-fine modeling of spatial-angular clues," in ECCV, 2018.
- [14] M. W. Tao, S. Hadap, J. Malik, and R. Ramamoorthi, "Depth from combining defocus and correspondence using light-field cameras," in *ICCV*, 2013, pp. 673–680.
- [15] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan, "Light field photography with a hand-held plenoptic camera," in *Stanford University Computer Science Tech Report*, 2005.

- [16] S. Wanner, S. Meister, and B. Goldlücke, "Datasets and benchmarks for densely sampled 4d light fields," in *Vision, Modeling & Visualization*, 2013, pp. 225–226.
- [17] T.-C. Wang, A. A. Efros, and R. Ramamoorthi, "Occlusion-aware depth estimation using light-field cameras," in *ICCV*, 2015, pp. 3487–3495.
- [18] H. Schilling, M. Diebold, C. Rother, and B. Jähne, "Trust your model: Light field depth estimation with inline occlusion handling," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4530–4538.
- [19] A. Alperovich, O. Johannsen, M. Strecke, and B. Goldluecke, "Light field intrinsics with a deep encoder-decoder network," in *CVPR*, 2018.
- [20] S. Wanner and B. Goldluecke, "Variational light field analysis for disparity estimation and super-resolution," *IEEE TPAMI*, vol. 36, no. 3, pp. 606–619, 2014.
- [21] Z. Zhang, Y. Liu, and Q. Dai, "Light field from micro-baseline image pair," in CVPR, 2015, pp. 3800–3809.
- [22] E. Penner and L. Zhang, "Soft 3D reconstruction for view synthesis," ACM TOG, vol. 36, no. 6, p. 235, 2017.
- [23] H. Zheng, M. Ji, H. Wang, Y. Liu, and L. Fang, "Crossnet: An endto-end reference-based super resolution network using cross-scale warping," in ECCV, 2018.
- [24] T. Zhou, R. Tucker, J. Flynn, G. Fyffe, and N. Snavely, "Stereo magnification: Learning view synthesis using multiplane images," ACM Transactions on Graphics (SIGGRAPH), 2018.
- [25] H. Shum and S. B. Kang, "Review of image-based rendering techniques," in *Visual Communications and Image Processing 2000*, vol. 4067. International Society for Optics and Photonics, 2000, pp. 2–13.
- [26] L. Shi, H. Hassanieh, A. Davis, D. Katabi, and F. Durand, "Light field reconstruction using sparsity in the continuous fourier domain," ACM TOG, vol. 34, no. 1, p. 12, 2014.
- [27] G. Wu, B. Masia, A. Jarabo, Y. Zhang, L. Wang, Q. Dai, T. Chai, and Y. Liu, "Light field image processing: An overview," *IEEE JSTSP*, vol. 11, no. 7, pp. 926–954, 2017.
- [28] R. Farrugia and C. Guillemot, "Light field super-resolution using a low-rank prior and deep convolutional neural networks," *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [29] W. Luo, Y. Li, R. Urtasun, and R. Zemel, "Understanding the effective receptive field in deep convolutional neural networks," in Advances in neural information processing systems, 2016, pp. 4898– 4906.
- [30] T. E. Bishop and P. Favaro, "The light field camera: Extended depth of field, aliasing, and superresolution," *IEEE TPAMI*, vol. 34, no. 5, pp. 972–986, 2012. [Online]. Available: http://dx.doi.org/10.1109/TPAMI.2011.168
- [31] G. Wu, Y. Liu, Q. Dai, and T. Chai, "Learning sheared epi structure for light field reconstruction," *IEEE Transactions on Image Processing*, vol. 28, no. 7, pp. 3261–3273, 2019.
- [32] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," arXiv preprint arXiv:1502.03167, 2015.
- [33] "Lytro." https://www.lytro.com/.
- [34] "ICME 2018 grand challenge on densely sampled light field reconstruction." http://www.tut.fi/civit/index.php/icme-2018-grandchallenge-densely-sampled-light-field-reconstruction/.
- [35] A. Dosovitskiy and T. Brox, "Generating images with perceptual similarity metrics based on deep networks," in Advances in neural information processing systems, 2016, pp. 658–666.
- [36] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for realtime style transfer and super-resolution," in *European conference on computer vision*. Springer, 2016, pp. 694–711.
- [37] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv preprint arXiv:1409.1556, 2014.
- [38] V. K. Adhikarla, M. Vinkler, D. Sumin, R. K. Mantiuk, K. Myszkowski, H.-P. Seidel, and P. Didyk, "Towards a quality metric for dense light fields," in *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, 2017, pp. 58–67.
- [39] H. G. Jeon, J. Park, G. Choe, and J. Park, "Accurate depth map estimation from a lenslet light field camera," in CVPR, 2015, pp. 1547–1555.
- [40] G. Chaurasia, S. Duchêne, O. Sorkine-Hornung, and G. Drettakis, "Depth synthesis and local warps for plausible image-based navigation," ACM TOG, vol. 32, 2013.

- [41] "Stanford (New) Light Field Archive." http://lightfield.stanford.edu/lfs.html.
- [42] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," Computer Science, 2014.
- [43] K. Yücer, A. Sorkine-Hornung, O. Wang, and O. Sorkine-Hornung, "Efficient 3D object segmentation from densely sampled light fields with applications to 3D reconstruction," ACM TOG, vol. 35, no. 3, 2016.
- [44] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous systems." http://tensorflow.org/.
- [45] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE TIP*, vol. 13, no. 4, pp. 600–612, 2004.
- [46] "Stanford Lytro Light Field Archive." http://lightfields.stanford.edu/.
- [47] Y. Wang, Y. Liu, W. Heidrich, and Q. Dai, "The light field attachment: Turning a dslr into a light field camera using a low budget camera ring," *IEEE transactions on visualization and computer* graphics, vol. 23, no. 10, pp. 2357–2364, 2016.
- [48] X. Lin, J. Wu, G. Zheng, and Q. Dai, "Camera array based light field microscopy," *Biomedical Optics Express*, vol. 6, no. 9, pp. 3179– 3189, 2015.
- [49] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [50] S. Iizuka, E. Simo-Serra, and H. Ishikawa, "Globally and locally consistent image completion," ACM Transactions on Graphics (ToG), vol. 36, no. 4, p. 107, 2017.



Gaochang Wu received the B.S. and M.S. degrees in mechanical engineering from Northeastern University, Shenyang, China, in 2013, and 2015, respectively. He is currently working toward Ph.D. degree in control theory and control engineering in Northeastern University, Shenyang, China. His current research interests include image processing, computational photography and machine learning.



Yebin Liu received the BE degree from Beijing University of Posts and Telecommunications, China, in 2002, and the PhD degree from the Automation Department, Tsinghua University, Beijing, China, in 2009. He has been working as a research fellow at the computer graphics group of the Max Planck Institute for Informatik, Germany, in 2010. He is currently an associate professor in Tsinghua University. His research areas include computer vision and computer graphics.



Lu FANG is currently an Associate Professor at Tsinghua University. She received her Ph.D in Electronic and Computer Engineering from HKUST in 2011, and B.E. from USTC in 2007, respectively. Dr. Fang's research interests include image / video processing, vision for intelligent robot, and computational photography. Dr. Fang serves as TC member in Multimedia Signal Processing Technical Committee (MMSP-TC) in IEEE Signal Processing Society.



Tianyou Chai received the Ph.D. degree in control theory and engineering from Northeastern University, Shenyang, China, in 1985. He has been with the Research Center of Automation, Northeastern University, Shenyang, China, since 1985, where he became a Professor in 1988 and a Chair Professor in 2004. His current research interests include adaptive control, intelligent decoupling control, integrated plant control and systems, and the development of control technologies with applications to various indus-

trial processes. Prof. Chai is a member of the Chinese Academy of Engineering, an academician of International Eurasian Academy of Sciences, IEEE Fellow and IFAC Fellow. He is a distinguished visiting fellow of The Royal Academy of Engineering (UK) and an Invitation Fellow of Japan Society for the Promotion of Science (JSPS).