



Light-field view synthesis using convolutional block attention module

Gul, M. Shahzeb Khan; Mukati, M. Umair; Bätz, Michel; Forchhammer, Søren; Keinert, Joachim

Published in:
Proceedings of 2021 IEEE International Conference on Image Processing

Link to article, DOI:
[10.1109/ICIP42928.2021.9506586](https://doi.org/10.1109/ICIP42928.2021.9506586)

Publication date:
2021

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Gul, M. S. K., Mukati, M. U., Bätz, M., Forchhammer, S., & Keinert, J. (2021). Light-field view synthesis using convolutional block attention module. In *Proceedings of 2021 IEEE International Conference on Image Processing* (pp. 3398-3402). IEEE. <https://doi.org/10.1109/ICIP42928.2021.9506586>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LIGHT-FIELD VIEW SYNTHESIS USING A CONVOLUTIONAL BLOCK ATTENTION MODULE

M. Shahzeb Khan Gul[†], M. Umair Mukati^{}, Michel Bätz[†], Søren Forchhammer^{*}, Joachim Keinert^{†*}*

[†] Moving Picture Technologies, Fraunhofer IIS, 91058, Erlangen, Germany

^{*} DTU Fotonik, Technical University of Denmark, Ørstedes Plads, Kgs. Lyngby - 2800, Denmark

ABSTRACT

Consumer light-field (LF) cameras suffer from a low or limited resolution because of the angular-spatial trade-off. To alleviate this drawback, we propose a novel learning-based approach utilizing attention mechanism to synthesize novel views of a light-field image using a sparse set of input views (i.e., 4 corner views) from a camera array. In the proposed method, we divide the process into three stages, stereo-feature extraction, disparity estimation, and final image refinement. We use three sequential convolutional neural networks for each stage. A residual convolutional block attention module (CBAM) is employed for final adaptive image refinement. Attention modules are helpful in learning and focusing more on the important features of the image and are thus sequentially applied in the channel and spatial dimensions. Experimental results show the robustness of the proposed method. Our proposed network outperforms the state-of-the-art learning-based light-field view synthesis methods on two challenging real-world datasets by 0.5 dB on average. Furthermore, we provide an ablation study to substantiate our findings.

Index Terms— Light-field, View synthesis, Deep-learning,

1. INTRODUCTION

Traditional cameras capture only the intensities of the incident light-rays at the photosensor, whereas, light-fields collect massive information of the scene by additionally acquiring the directional information of the incident light-rays. This additional angular information becomes beneficial in many applications such as depth estimation, post-capture refocusing, and 3D reconstruction.

Light-fields can be captured through many different ways. Conventional methods include camera arrays [1], which can capture a LF simultaneously in time but have a limited angular resolution because of the numbers of cameras and the necessary spacing. Gantry systems [2], which sample a high angular resolution light-field by sequentially capturing each viewpoint, cannot capture dynamic content. On the other hand, hand-held light-field cameras, such as from Lytro and Raytrix [3], offer a feasible solution for light-field capture, based on the plenoptic camera designs [4]. However, it suffers from the spatial and angular resolution trade-off due to the sensor limitation.

*Published in the IEEE 2021 International Conference on Image Processing (IEEE ICIP 2021), scheduled for 19-22 September 2021 in Anchorage, Alaska, United States. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works, must be obtained from the IEEE. Contact: Manager, Copyrights and Permissions / IEEE Service Center / 445 Hoes Lane / P.O. Box 1331 / Piscataway, NJ 08855-1331, USA. Telephone: + Intl. 908-562-3966.

In recent years, different light-field view synthesis approaches have been published aiming to solve the limited angular resolution problem. The success of deep learning compared to traditional methods for numerous computer vision and image processing problems, such as multi-view stereo [5,6], optical flow [7], and super-resolution [8], inspired the researchers to develop learning-based view interpolation methods [9–14]. These methods can be broadly divided into two categories: image-based rendering (IBR) and depth image-based rendering (DIBR) methods. The IBR methods synthesizing the novel views without explicitly modeling the scene depth [13, 14], tend to fail with increasing baseline of the input views, whereas, depth-based methods can handle inputs with wider baseline (i.e. large disparities) [11, 15].

In this paper, we demonstrate for the first time that CBAM used in a different context, i.e., object detection, can be used for light-field view synthesis. We divide the whole process into three stages, i.e., stereo-feature extraction, disparity estimation, and final image refinement. In the first step, we extract features directly from stereo pairs and we feed these stereo features into a disparity estimation network to estimate target disparity maps corresponding each input view. Multiple disparity vectors per target view alleviate the interpolation error due to the warping of input images. In the end, an image refinement network using a convolutional block attention module takes the warped images as input, and outputs the final novel target view. Attention modules have been studied and used to a great extent in the literature [16], [17], [18]. Attention mechanisms not only guide the network to focus on important features, but they also improve the representation of the regions of interest.

The remainder of the paper is structured as follows. In Section 2, we provide a comprehensive literature review on state-of-the-art view synthesis algorithms. Section 3 explains the proposed network architecture. Experimental results are then presented and discussed in Section 4 before the paper is concluded in Section 5.

2. RELATED WORK

2.1. Image-based rendering (IBR) method

In literature, some methods approach the light-field view synthesis problem as a specific case of signal reconstruction. They exploit the intrinsic sparsity of the light-field images. In [19], a captured light-field is considered to be sparse in the continuous Fourier domain; hence, novel views can be reconstructed from a small number of 1D viewpoint trajectories. Levin and Durand [20] synthesize the novel views using 3D focal stack sequences with an assumption that the light-field contains only Lambertian surfaces.

Many researchers used sparse coding as a tool to interpolate novel views. In [21], a global dictionary is learned from light-field patches to synthesize novel views from a set of input light-field




Fig. 1: Overall flow of the proposed network. Here, I_L , I_R , and I_B are the selected input images as shown in Fig. 2. U and V are of the size of the input views and repeatedly contain the u and v coordinate of the target view to be rendered, respectively.

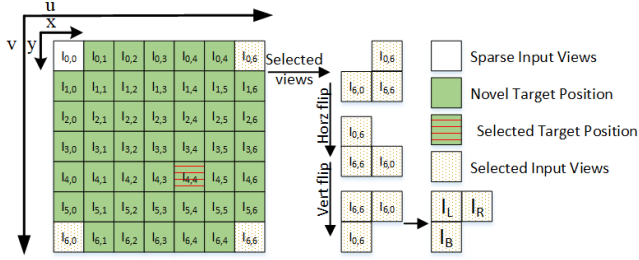


Fig. 2: Based on the angular coordinate of the target view, we select three input views from a given sparse set of four corner views. Input views with the least absolute distance from the target view are selected. To maintain the correspondence direction constant among different input combinations, we flip the selected inputs views in horizontal or vertical direction accordingly to make an upper left triangle shape.

views. On the other hand, a local dictionary learned using the central region of the light-field images is successfully applied for the interpolation of novel views in [22], [23].

In [24] and [10], the authors approached the light-field view synthesis problem as angular domain super-resolution. In [24], two sequential convolutional neural networks (CNNs) take sub-aperture views as inputs to increase the spatial and angular resolution of the light-field consecutively. On the other hand, [10] applied the CNNs directly on lenslet images to increase the resolution in both spatial and angular domains. Wang *et al.* [14] proposed to use stacked 3D volumes of epipolar plane images in a CNN framework. In [13], spatio-angular alternating convolutions (i.e., pseudo 4D-CNN) are performed to incorporate 4D light-field data for novel view interpolation. On the other hand, our proposed network synthesizes the novel view by explicitly modeling the scene geometry.

2.2. Depth image-based rendering (DIBR) method

Typically, depth-based view synthesis methods are divided into two steps. At first, these methods estimate the disparity at a target view using the sparse light-field views and then utilize it to synthesize the target view by warping of the input views. In [9], two CNNs are trained to jointly estimate the depth and synthesize a view in an end-to-end fashion. In [11], a feature extraction CNN is introduced as the first part of the network architecture. Moreover, the methods estimates four disparity maps corresponding to each input view and then utilize a selection CNN for the selection of pixels from warped images. Jin *et al.* [25] proposed CNN estimating 4D depth maps for a high angular resolution light-field, i.e., providing depth for each light ray in the 4D light-field. The resulting 4D depth is then utilized

to synthesize all novel views by backward warping. As a final step, a CNN light-field blending module is employed to get the final light-field.

Recently, a view synthesis method based on multi-plane image (MPI) [26] was introduced, where a scene is represented by a stack of RGBA planes at different depths. In [27], this idea is extended to light-field using a 3D-CNN to infer an MPI. Flynn *et al.* [28] proposed to use the variational optimization framework in conjunction with deep learning to improve the reconstruction quality of MPIs.

While our algorithm is inspired by [11], it introduces a new combination of attention mechanism to efficiently deal with the occluded regions as shown in Fig. 4.

2.3. Attention mechanism

Very recently, attention mechanisms have become a very popular method to enhance a deep neural network. Attention modules try to make a neural network focus on salient regions of its feature representation. In [16], a residual attention network is used for image classification. The Squeeze-and-excitation (SE) block attention module introduced in [17], improves the classification accuracy by exploiting channel inter-dependencies. In [18], the authors introduced a CBAM by adding a spatial attention mechanism to a SE block to validate performance of object detection. In [29] a variant of CBAM is used in video frame interpolation to refine the interpolated frame. Our proposed method uses the attention mechanism for a similar purpose, i.e., image refinement. We integrate CBAM as proposed in [18] into our light-field view interpolation pipeline, using both channel and spatial attention.

3. PROPOSED METHOD

Let $L(x, y, u, v)$ denote a light-field according to two-plane parametrization as shown in Fig. 2, where (x, y) are the spatial coordinates and (u, v) are the angular coordinates. Each sub-aperture view in the light-field is denoted as $I_{u,v}$.

3.1. Network Architecture

The architecture of our proposed network is shown in Fig. 1. To synthesize high-quality novel views, we divide the proposed method into three stages: stereo feature extraction, disparity estimation, and image refinement. Let L^S be four corner views of a light-field, our proposed method makes use of the three closest corner views to reconstruct a high-quality image I_{u_t, v_t} at the target position:

$$\hat{I}_{u_t, v_t} = f(I_L, I_R, I_B, u_t, v_t), \quad (1)$$

where f is the function that we aim to learn and which should estimate a view \hat{I}_{u_t, v_t} as similar as possible to the ground-truth image

I_{u_t, v_t} , I_L , I_R , and I_B are the three selected corner views for the target position (u_t, v_t) . Out of the four, the three reference views are selected as input to the network based on their minimum distance to the target view (as depicted in Fig. 2). Selecting three corner views as opposed to four reduces the overall complexity of the network. To maintain constant geometry among different inputs, we flip the images left-to-right and/or up-to-down accordingly so that we always have an upper left triangle shape of the selected input views (see Fig. 2). Note that the flipping operation will also change the angular coordinate of the target view. Similarly, we flip back the output of the network to get the final reconstructed image.

3.1.1. Stereo feature extraction network

We first extract stereo features from the sparse set of selected input views. The architecture of the feature extraction module f_e is based on [11]. The network f_e consists of six convolutional layers with 3×3 kernels and two average pooling layers with 16×16 and 8×8 kernels. However, instead of extracting features for each image individually, as done in [11], we extract features directly from stereo pairs. We group the input views into three stereo pairs, namely horizontal (I_L, I_R), vertical (I_L, I_B), and diagonal (I_B, I_R). Horizontal and vertical stereo pairs use the same architecture for feature extraction. The reason for having a shared network here is that the geometry between the images can be made identical to the horizontal case when the vertical stereo pair is rotated by 90° counter-clockwise. Furthermore, this reduces the complexity of the network and thus simplifies the trainability. In addition to the stereo pair as input, the network also takes U and/or V image planes, where $U(x, y) = u_t, V(x, y) = v_t \quad \forall (x, y)$. These image planes aid the network in collecting appropriate information from each stereo pair based on the position of the novel view.

Let $F_{\text{horz}} = f_e(I_L, I_R, U)$, $F_{\text{vert}} = f_e(I_L, I_B, V)$, and $F_{\text{diag}} = f_e(I_B, I_R, U, V)$ be the computed feature volumes for the three stereo pairs. These 32-channel feature volumes are then concatenated and passed to the disparity estimation network.

3.1.2. Disparity estimation network

The goal of this module is to estimate a disparity map per input view valid at the target position. Multiple disparity values per target view pixel allow for a more effective handling of occluded regions as compared to [9] where a single disparity map is estimated for all input views. The architecture of the disparity estimation network f_d is inspired by [11]. The network consists of seven convolutional layers with 3×3 filter size. The first four layers use dilated convolutions at rates 2, 4, 8, and 16, respectively. The network takes the features extracted from the stereo pairs and angular coordinate image planes U and V as input and output disparity maps D_i , for $i \in \{L, R, B\}$.

3.1.3. Image refinement network

The estimated disparity maps are used to warp each corner view in order to have them registered with the target one. After warping the input views I_i^w , for $i \in \{L, R, B\}$, at the target location using the estimated disparity maps, we estimate the final image using a residual CBAM refinement network. The architecture of the refinement network consists of a head and tail 3×3 convolution layer with five residual groups in between, referred to as *ResGroups* from now, each of which consists of three CBAMs [18]. As shown in Fig. 3 (a), each CBAM has two 3×3 convolution layers with ReLU activation in between, followed by a channel attention (CA) and spatial attention

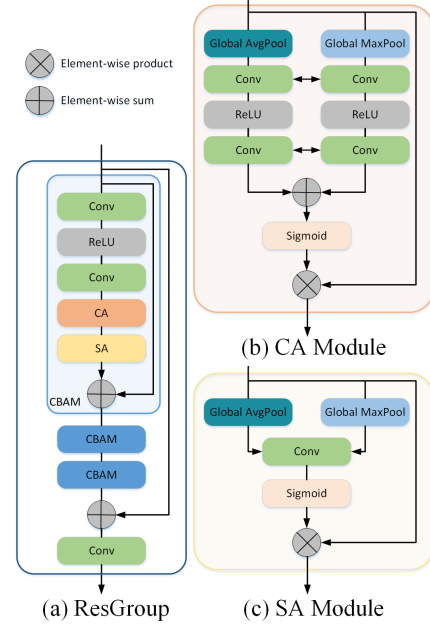


Fig. 3: Illustration of ResGroup, CA module, and SA module.

(SA) module before the residual connection. Both attention modules, channel and spatial, compute complementary attentions focusing on ‘what’ and ‘where’ respectively. The architecture of the CA module is shown in Fig. 3 (c). We first aggregate the spatial information using average-pooling and max-pooling operations. Features descriptors from both operations are then forwarded to a shared two 1×1 convolution layer network to produce a 1D channel attention map. Similarly, the architecture of the SA module is illustrated in Fig. 3 (b). We aggregate the channel information using two pooling operations followed by a 7×7 convolution layer producing a 2D spatial attention map. The final output image at the target position is obtained by adding the output of the refinement network and the average of the warped views.

3.2. Loss Function

The model is trained by optimizing the sum of the following loss functions:

$$\mathcal{L}_{\text{final}} = \|I_{u,v} - \hat{I}_{u,v}\|_1 + \lambda_1 \|\nabla I_{u,v} - \nabla \hat{I}_{u,v}\|_1, \quad (2)$$

$$\mathcal{L}_{\text{warp}} = \lambda_2 \sum_i \|I_{u,v} - I_i^w\|_1 + \lambda_3 \sum_i \|\nabla I_{u,v} - \nabla I_i^w\|_1. \quad (3)$$

I_i^w is the warped view from one stereo pair using the estimated disparity map, where $i \in \{L, R, B\}$. $I_{u,v}$ and $\hat{I}_{u,v}$ are ground-truth and estimated novel views, respectively. Moreover, we experimentally set $\lambda_1 = 0.5$, $\lambda_2 = 0.25$, and $\lambda_3 = 0.125$. The first term in both losses is the L_1 loss between the predicted and the ground-truth images, whereas, the second term preserves texture by taking L_1 -norm between spatial gradients.

3.3. Training

The proposed model is implemented using TensorFlow as framework. The model is trained on the light-field dataset *Flowers* presented in [30]. This dataset consists of 3343 images of flowers. We

Method	<i>Diverse</i> [9]		<i>Flowers</i> [30]	
	PSNR (dB)	MS-SSIM (dB)	PSNR (dB)	MS-SSIM (dB)
Kalantari <i>et al.</i> [9]	39.56	24.31	40.22	22.47
Navarro <i>et al.</i> [11]	40.75	25.57	42.67	25.04
LFVS (ours)	41.14	25.89	43.13	25.44
LFVS-AM (ours)	41.33	26.05	43.17	25.54

Table 1: Quantitative results on the test datasets *Flowers* and *Diverse*. For each row, average PSNR (dB) and MS-SSIM (dB) values for Y channel of the image achieved on all intermediate rendering positions are listed for all methods.




Fig. 4: Visual comparison of the central view of the scenes. First and second row contains the LEAVES scene and its zoomed-in region from *Diverse* dataset, whereas, third and fourth row contains the IMG_7824 scene and its zoomed-in region from *Flowers* dataset. (a)(f) Kalantari *et al.* [9], (b)(g) Navarro *et al.* [11], (c)(h) LFVS, (d)(i) LFVS-AM, and (e)(j) Ground-truth.

randomly select 100 images for testing and used the remaining images for training the network. Each light-field in the dataset has a spatial resolution of 540×372 and an angular resolution 14×14 . During the experiment, we only consider the center 7×7 grid of views because the corner views are affected by vignetting. As a data augmentation, we first randomly crop inputs to a size of 192×192 then randomly perform gamma correction using a value from the range $[0.4, 1]$. The pixel values are normalized to a range of $[-1, 1]$. The optimization of the proposed network is done with the ADAM optimizer where β_1 and β_2 are set to 0.9 and 0.99. The learning rate is set to 0.0001, while the batch size is 8. The model converges after 1000 epochs, and it takes approximately 2.5 days to train on a GeForce GTX 1080 Ti GPU.

4. EXPERIMENTAL RESULTS

In this section, we compare the proposed method with state-of-the-art methods both quantitatively and qualitatively. We evaluate all the methods on two test datasets, one contains 100 light-fields of flowers that we randomly select from [30]. The other test set *Diverse* consists of 30 light-fields [9] – mostly outdoor scenes. Besides, we also conducted ablation studies to assess the effectiveness of the various components of our network.

4.1. Comparison with the state-of-the-art

To evaluate the performance of the proposed method, named LFVS-AM (Light-Field View Synthesis using Attention Module), we compare it with state-of-the-art methods including Kalantari *et al.* [9]

# ResGroups	# CBAMs	PSNR (dB)	MS-SSIM (dB)
3	3	41.16	25.76
3	5	41.25	25.95
5	3	41.33	26.05

Table 2: Quantitative results for three different models with varying number of *ResGroups* and CBAMs on *Diverse* [9] dataset.

and Navarro *et al.* [11]. For quantitative evaluation, we calculate the PSNR and MS-SSIM (dB) = $-10 \log_{10}(1 - \text{MS-SSIM})$ values for the luminance channel of the images between the estimation and the ground-truth. Table 1 shows the PSNR and MS-SSIM results averaged over all intermediate target positions – 7×7 views excluding the 4 corner views which served as input. It is evident that the proposed method produces better results compared to both state-of-the-art methods across all metrics on both the datasets. Our method outperforms [9] by 2.3 dB and [11] by 0.5 dB on average. Moreover, visual results in Fig. 4 demonstrates that Kalantari *et al.* [9] and Navarro *et al.* [11] failed to generate plausible results in the presence of occlusions. For example in Fig. 4a and 4b, both of these methods are not able to generate the pole between the two leaves, which is occluded in the input views. In comparison, our method better estimates the occluded region which is closer to the ground-truth, as shown in Fig. 4d.

4.2. Ablation Study

4.2.1. Effects of hyperparameters

In Table 2, we compare three variants of LFVS-AM models having a different number of *ResGroups* and CBAMs. The metrics in the table are calculated for Y channel of the image. From the results, we can conclude that the model with 5 *ResGroup* and 3 CBAM performs the best among the three models, but the differences in performance are not that significant when the total number of CBAM blocks becomes the same i.e., $(\# \text{ResGroups}) \times (\# \text{CBAMs}) = 15$.

4.2.2. Effect of the attention mechanism

To evaluate the effect of the attention mechanism in the refinement network, we compare the proposed method with the model having no attention mechanism. For the model without attention mechanism (LFVS), we remove the CA and SA modules in Fig. 3 (a) to change the CBAM to a residual block. In Table 1, average PSNR and MS-SSIM (dB) values on all intermediate rendering positions are presented. The results indicate the effectiveness of the attention mechanism in our proposed refinement network. Our LFVS model without attention mechanism is better than SoTA because of the residual connections in the network but still can not handle occlusions, as shown in Fig. 4c. In contrast, the attention module guides the network to efficiently predict the occluded regions, see Fig. 4d.

5. CONCLUSION

In this work, we proposed a novel end-to-end CBAM light-field view synthesis method to reconstruct dense light-fields from sparse input data. Our model consists of stereo feature extraction, disparity estimation, and residual CBAM image refinement networks. All these components are modelled using three sequential CNNs. We presented quantitative and qualitative results of the proposed method on two different datasets. Experimental results show that our model outperforms state-of-the-art approaches by 0.5 dB on average. Future

work will include to analyze the potential of our proposed method in the context of light-field compression frameworks.

6. ACKNOWLEDGMENT

We thank Milan Stepanov (L2S, CentraleSupélec) for his valuable help during the re-implementation of [11]. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 765911 (RealVision).

7. REFERENCES

- [1] Bennett Wilburn, Neel Joshi, Vaibhav Vaish, Eino-Ville Talvala, Emilio Antunez, Adam Barth, Andrew Adams, Mark Horowitz, and Marc Levoy, “High performance imaging using large camera arrays,” in *Proc. of ACM Trans. on Graphics*, pp. 765–776, July 2005, Los Angeles, California.
- [2] MSK Gul, Thorsten Wolf, Michel Bätz, Matthias Ziegler, and Joachim Keinert, “A high-resolution high dynamic range light-field dataset with an application to view synthesis and tone-mapping,” in *Proc. of IEEE Int. Conf. on Multimedia and Expo Workshops*, July, 2020, pp. 1–6, London, UK.
- [3] “Raytrix: Light filed technology,” 2018, <https://raytrix.de/technology/>.
- [4] Ren Ng, Marc Levoy, Mathieu Brédif, Gene Duval, Mark Horowitz, and Pat Hanrahan, *Light field photography with a hand-held plenoptic camera*, Ph.D. thesis, Stanford University, 2005.
- [5] Matteo Poggi, Davide Pallotti, Fabio Tosi, and Stefano Mattoccia, “Guided stereo matching,” in *Proc. of IEEE Int. Conf. on Computer Vision and Pattern Recognition*, June 2019, pp. 979–988, Long Beach, USA.
- [6] MSK Gul, Michel Bätz, and Joachim Keinert, “Pixel-wise confidences for stereo disparities using recurrent neural networks,” in *British Machine Vision Conf.*, Sep. 2019, p. 23, Cardiff, UK.
- [7] Pengpeng Liu, Michael Lyu, Irwin King, and Jia Xu, “Selfflow: Self-supervised learning of optical flow,” in *Proc. of IEEE Int. Conf. on Computer Vision and Pattern Recognition*, June 2019, pp. 4571–4580, Long Beach, USA.
- [8] Thomas Köhler, Michel Bätz, Farzad Naderi, André Kaup, Andreas Maier, and Christian Riess, “Toward bridging the simulated-to-real gap: Benchmarking super-resolution on real data,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 42, no. 11, pp. 2944–2959, 2020.
- [9] Nima Khademi Kalantari, Ting-Chun Wang, and Ravi Ramamoorthi, “Learning-based view synthesis for light field cameras,” in *ACM Trans. on Graphics*, July 2016, pp. 1–10, Anaheim, USA.
- [10] MSK Gul and Bahadır K Gunturk, “Spatial and angular resolution enhancement of light fields using convolutional neural networks,” *IEEE Trans. on Image Processing*, vol. 27, no. 5, pp. 2146–2159, 2018.
- [11] Julia Navarro and Neus Sabater, “Learning occlusion-aware view synthesis for light fields,” *arXiv preprint arXiv:1905.11271*, 2019.
- [12] Lu Liu, Zicheng Nian, and Cheolkon Jung, “DCM-CNN: Densely connected multiloss convolutional neural networks for light field view synthesis,” *IEEE Access*, vol. 8, pp. 78542–78552, 2020.
- [13] Henry Wing Fung Yeung, Junhui Hou, Jie Chen, Yuk Ying Chung, and Xiaoming Chen, “Fast light field reconstruction with deep coarse-to-fine modeling of spatial-angular clues,” in *Proc. of the European Conf. on Computer Vision*, Sep. 2018, pp. 137–152, Munich, Germany.
- [14] Yunlong Wang, Fei Liu, Zilei Wang, Guangqi Hou, Zhenan Sun, and Tieniu Tan, “End-to-end view synthesis for light field imaging with pseudo 4DCNN,” in *Proc. of the European Conf. on Computer Vision*, Sep. 2018, pp. 333–348, Munich, Germany.
- [15] Dingcheng Yue, MSK Gul, Michel Bätz, Joachim Keinert, and Rafat Mantiuk, “A benchmark of light field view interpolation methods,” in *Proc. of IEEE Int. Conf. on Multimedia and Expo Workshops*, July, 2020, pp. 1–6, London, UK.
- [16] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Hong-gang Zhang, Xiaogang Wang, and Xiaoou Tang, “Residual attention network for image classification,” in *Proc. of IEEE Int. Conf. Computer Vision and Pattern Recognition*, July, 2017, pp. 3156–3164, Hawaii, USA.
- [17] Jie Hu, Li Shen, and Gang Sun, “Squeeze-and-excitation networks,” in *Proc. of IEEE Int. Conf. on Computer Vision and Pattern Recognition*, July, 2018, pp. 7132–7141, Utah, USA.
- [18] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon, “CBAM: Convolutional block attention module,” in *Proc. of the European Conf. on Computer Vision*, Sep. 2018, pp. 3–19, Munich, Germany.
- [19] Lixin Shi, Haitham Hassanieh, Abe Davis, Dina Katabi, and Fredo Durand, “Light field reconstruction using sparsity in the continuous fourier domain,” *ACM Trans. on Graphics*, pp. 1–13, Aug. 2014, Vancouver, USA.
- [20] Anat Levin and Fredo Durand, “Linear view synthesis using a dimensionality gap light field prior,” in *Proc. of IEEE Int. Conf. on Computer Vision and Pattern Recognition*, June, 2010, pp. 1831–1838, San Francisco, USA.
- [21] Kshitij Marwah, Gordon Wetzstein, Yosuke Bando, and Ramesh Raskar, “Compressive light field photography using overcomplete dictionaries and optimized projections,” *ACM Trans. on Graphics*, pp. 1–12, July 2013, Anaheim, USA.
- [22] David C Schedl, Clemens Birklbauer, and Oliver Bimber, “Directional super-resolution by means of coded sampling and guided upsampling,” in *Proc. of IEEE Int. Conf. on Computational Photography*, April 2015, pp. 1–10, Houston, USA.
- [23] David C Schedl, Clemens Birklbauer, and Oliver Bimber, “Optimized sampling for view interpolation in light fields using local dictionaries,” *Elsevier Computer Vision and Image Understanding*, pp. 93–103, 2018.
- [24] Youngjin Yoon, Hae-Gon Jeon, Donggeun Yoo, Joon-Young Lee, and In So Kweon, “Light-field image super-resolution using convolutional neural network,” *IEEE Signal Processing Letters*, vol. 24, no. 6, pp. 848–852, 2017.
- [25] Jing Jin, Junhui Hou, Hui Yuan, and Sam Kwong, “Learning light field angular super-resolution via a geometry-aware network,” in *AAAI*, Feb. 2020, pp. 11141–11148, New York, USA.
- [26] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely, “Stereo magnification: learning view synthesis using multi-plane images,” *ACM Trans. on Graphics*, Aug. 2018, Vancouver, USA.
- [27] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar, “Local light field fusion: Practical view synthesis with prescriptive sampling guidelines,” *ACM Trans. on Graphics*, Jul. 2019, Los Angeles, USA.
- [28] John Flynn, Michael Broxton, Paul Debevec, Matthew DuVall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker, “Deep-view: View synthesis with learned gradient descent,” in *Proc. of IEEE Int. Conf. on Computer Vision and Pattern Recognition*, June 2019, pp. 2367–2376, Long Beach, USA.
- [29] Keito Suzuki and Masaaki Ikehara, “Residual learning of video frame interpolation using convolutional LSTM,” *IEEE Access*, vol. 8, pp. 134185–134193, 2020.
- [30] Pratul P Srinivasan, Tongzhou Wang, Ashwin Sreelal, Ravi Ramamoorthi, and Ren Ng, “Learning to synthesize a 4d RGBD light field from a single image,” in *Proc. of IEEE Int. Conf. Computer Vision and Pattern Recognition*, July, 2017, pp. 2243–2251, Hawaii, USA.