

# Deep Sparse Light Field Refocusing

Shachar Ben Dayan  
bendayan@mail.tau.ac.il

David Mendlovic  
mend@eng.tau.ac.il

Raja Giryes  
raja@tauex.tau.ac.il

School of Electrical Engineering  
The Iby and Aladar Fleischman Faculty  
of Engineering  
Tel Aviv University  
Tel Aviv, Israel

---

## Abstract

Light field photography enables to record 4D images, containing angular information alongside spatial information of the scene. One of the important applications of light field imaging is post-capture refocusing. Current methods require for this purpose a dense field of angle views; those can be acquired with a micro-lens system or with a compressive system. Both techniques have major drawbacks to consider, including bulky structures and angular-spatial resolution trade-off. We present a novel implementation of digital refocusing based on sparse angular information using neural networks. This allows recording high spatial resolution in favor of the angular resolution, thus, enabling to design compact and simple devices with improved hardware as well as better performance of compressive systems. We use a novel convolutional neural network whose relatively small structure enables fast reconstruction with low memory consumption. Moreover, it allows handling without re-training various refocusing ranges and noise levels. Results show major improvement compared to existing methods.

## 1 Introduction

Light field photography has attracted significant attention in recent years due to its unique capability to extract depth without active components [1, 2, 3]. While 2D cameras only capture the total amount of light at each pixel on the sensor, namely, the projection of the light in the scene, light field cameras also record the direction of each ray intersecting with the sensor in a single capture. Thus, light field images contain spatial and angular information of a given scene. A light field image can be represented as a 4D tensor, carrying the dimensions of the RGB color channels and the multiple angle views alongside the spatial dimensions.

Early methods for light field capturing include multi 2D-cameras array [4], where each camera captures the scene from a different angle; and sequential imaging [5], where a single camera is shifted and takes multiple images. Both strategies are bulky and not suitable for capturing dynamic scenes. Lytro, a plenoptic camera [6] captures a light field image in a single exposure. The device is equipped with a micro-lens array located between the sensor and the main lens. This system enables the recording of the direction of the rays arriving at the sensor plane. However, since the sensor is meant to capture both spatial and angular information, there is an inherent tradeoff between the captured spatial and angular resolution.

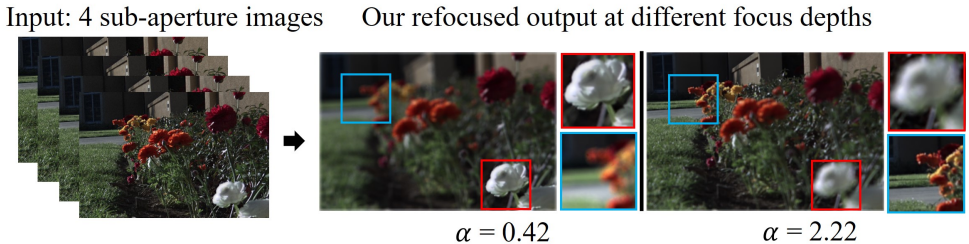


Figure 1: Our proposed method of light field refocusing needs only 4 sub-aperture input views to output a good refocused image at a specific depth.

Other methods such as coded masks [58] or coded aperture [20] based light field also suffer from the above trade-off and produce images with low spatial resolution. Recent approaches suggest using compressive systems for light field acquisition [6, 2, 23, 26], where the camera records a coded projection of the light rays intersecting with the sensor. The light field images are then recovered from their projections using methods such as sparse coding and deep learning. However, the more angular information there is to compress, the more artifacts the output images have, resulting in poor resolution and signal-to-noise ratio (SNR).

**Contribution.** Given the above, it is important to perform light field applications such as refocusing, which is the focus of this work, relying only on sparse angular data. We propose a technique for post-capture refocusing based on sparse light field data: only  $2 \times 2$  sub-aperture views. To this end, we exploit the power of convolutional neural networks (CNNs). The four input views are fed into a CNN, which outputs a refocused image according to a given focus plane (see Fig. 1). To the best of our knowledge, we are the first to suggest an end-to-end light field refocusing based on deep learning. Our approach shows improved performance in terms of PSNR, SSIM and computational time, compared to the current existing methods, which are of two kinds: direct refocusing, and indirect synthesis methods that synthesize from a small number of views a full light field to generate from it the refocused image. Our trained network, RefocusNet, exhibits very good generalization performance across different sensors. We demonstrate that after training it on a dataset from one sensor, it can be used with another different light field system without the need of a new training.

## 2 Related Work

While there are many works on refocusing [9, 22, 29, 36, 39], this work focuses on light field based refocusing. The basic light field refocusing algorithm relies on shifting the sub-aperture views according to the desired focus plane and the disparity of each sub-aperture view and then averaging the shifted images [20, 37]. This method has been demonstrated on light field images captured by a Lytro camera [28] and was recently generalized to tilt-shift photography [3]. However, this method requires a dense angular sampled light field, or otherwise the resulted images will contain artifacts (see Section 4). Several methods locate and fix the aliasing caused by the lack of angular information [18, 43].

Another body of works uses the Fourier Slice Theorem [6, 7, 22]. The Fourier Slice Digital Refocusing method [27] proves that spatial light field refocusing is equal to computing the Fourier transform of a 4D light field, selecting a relevant 2D slice and then performing

2D inverse Fourier transform. Yet, it does not perform well on sparse angular data. A recent strategy [14] represents light fields efficiently using Fourier Disparity Layers (FDL). This compact representation can be used for sub-aperture views rendering and direct refocusing.

Other methods suggest increasing the angular resolution of a light field by reconstructing additional sub-aperture views and produce the refocused images by shifting and averaging the sub-aperture images [15]. With the grown popularity of deep learning methods in computer vision tasks, a CNN has been trained to reconstruct a  $7 \times 7$  light field from  $3 \times 3$  given angles, while using EPIs (epipolar plane images) [16]. While this method outperforms state-of-the-art methods, they indicate that the algorithm will not perform well on inputs with less than  $3 \times 3$  angle views. The work of Kalantari *et al.* [17] uses two sequential CNNs to synthesize an  $8 \times 8$  light field from 4 ( $2 \times 2$ ) corner sub-aperture images. Although this method performs well on sparse angular data, the fact that each view needs to be created separately is time-consuming, which makes the algorithm not suitable for real-time applications. This is a major drawback when relying on light field synthesis methods [11, 13, 14, 16] for refocusing; instead, direct refocusing can be performed using the given sub-aperture views, similar to image fusion networks [18, 19].

### 3 Mathematical Background

**Light Field Photography.** Levoy *et al.* [19] introduced the modeling of a 4D light field as a two-plane parameterization, which simplifies the problem formulation. We briefly describe it here. The common light field device is usually equipped with a micro-lens array located between the sensor and the main lens, which is placed at the aperture plane. Given the sensor plane  $(x, y)$ , the aperture plane  $(u, v)$ , and the distance between them  $F$  (focal length), we may say that the intensity of each point  $(x, y)$  on the sensor plane equals to the integration of all the light rays moving through the whole aperture reaching this point:

$$E_F(x, y) = \frac{1}{F^2} \iint L_F(x, y, u, v) A(u, v) \cos^4 \phi \, du \, dv. \quad (1)$$

$E_F(x, y)$  is the intensity at point  $(x, y)$  on the sensor plane (known as image plane),  $L_F$  is the radiance along the ray from  $(u, v)$  to  $(x, y)$ ,  $A$  is the aperture function (1 inside the aperture and 0 outside), and  $\phi$  is the angle between the image plane normal and the light ray. Let

$$\overline{L}_F(x, y, u, v) = L_F(x, y, u, v) \cos^4 \phi, \quad (2)$$

then if we assume that  $E_F$  is zero outside the aperture boundaries, we get the simplified form:

$$E_F(x, y) = \frac{1}{F^2} \iint \overline{L}_F(x, y, u, v) \, du \, dv. \quad (3)$$

**Digital Refocusing.** An important light field application is post-capture refocusing. As 2D images only contain spatial information, we have to choose the focus plane before capturing the image and adjust the aperture accordingly. In light field photography, the angular information is also available, which allows us to change the focus plane also after capturing.

To produce a refocused light field image, one has to artificially change the current sensor plane. As all possible focus planes are parallel, we need to multiply the current plane, located at distance  $F$  from the aperture, by a factor  $\alpha$  to get the new focus plane,  $F'$ . Fig. 2 shows

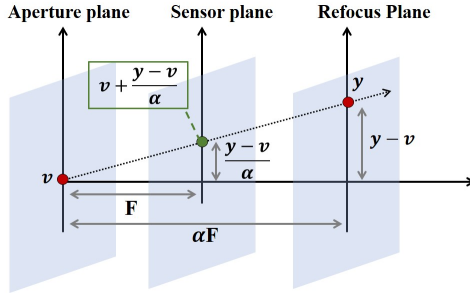


Figure 2: Changing the sensor plane artificially to get a refocused image. The new plane is located at a distance of  $F' = \alpha F$  from the aperture plane.

that using triangles similarity, we can compute the new intersection point of the ray with the sensor plane. Substituting  $(x, y)$  with the new coordinates and  $F' = \alpha F$  in (3), we get:

$$E_{(\alpha F)}(x, y) = \frac{1}{\alpha^2 F^2} \iint \overline{L}_F(u + \frac{x-u}{\alpha}, v + \frac{y-v}{\alpha}, u, v) dudv. \quad (4)$$

From (4) we may conclude that practically, creating a refocused image is done by shifting the sub-aperture images, each by a different factor determined by the focus factor and view location on the rectangular grid, followed by averaging the shifted images (each corresponding to a different sub-aperture view). This can be explained from an optical point of view: each sub-aperture image represents a different pinhole on the aperture plane. Since light field is the summation of rays intersecting at a specific pixel on the sensor, we sum the sub-aperture images. In order to embody the location of each pinhole on the aperture plane in the refocused image, we shift the sub-aperture images.

The refocused image model in Eq. (4) is continuous. However, a single light field image is a discrete set of 2D images. In order to imitate this continuity and create a smooth image, the light field set needs to be as dense as possible, i.e., be with a high angular resolution. On the other hand, because of the spatial-angular resolution trade-off mentioned above, we would like to be able to decrease the angular resolution as much as possible and still get a high quality refocused image. Our proposed method, described next, targets this challenge.

## 4 Proposed Method

Our goal is to address the spatial-angular resolution trade-off and create a light field refocusing program that operates in real-time with low computational complexity. Thus, we propose a convolutional neural network (CNN) based approach for refocusing from sparse light field. **Sparse Angular Data Refocusing.** Based on the assumption that we only have 4 input views arranged on a  $2 \times 2$  grid, we first check what happens when we refocus an image only from these views, without any priors. We shift each view according to its location on the grid and the target focus factor, and then sum the shifted views and divide them by 4. We use the code in [2]. As expected, the resulted image has noticeable ghosting artifacts around objects that are not on the selected focus plane (see Fig.3 (left)).

Note that the differences around the region in focus (the central flower) in Fig. 3, between the refocusing based on  $2 \times 2$  views and the refocusing based on a higher number of views



Figure 3: Demonstration of conventional light field refocusing based on a varying number of sub-aperture views. The ghosting artifacts are clearly seen in the 2x2 grid result. Also, minor artifacts are still seen in the 9x9 grid result (note the tree and the edge of the car). However, we can barely see any artifacts in the 17x17 grid.

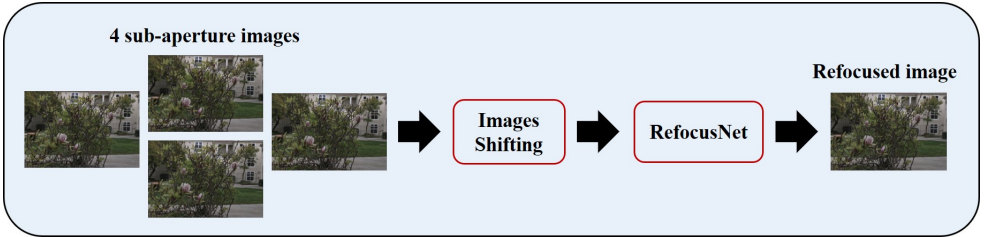


Figure 4: Flow chart of the proposed algorithm. Given 4 sub-aperture images, we shift the views according to their locations on the rectangular light field grid and the randomly chosen focus factor. Then we feed them into a CNN that outputs a refocused RGB image.

are negligible; the reason is that when we refer to an object at the focus plane, the shift of the sub-aperture images equals to zero, and all the shifted sub-aperture images are the same and averaging them results in the same image, making the object appear sharp and clear. On the contrary, when looking at objects outside of the focus plane, we expect them to appear blurred. Theoretically, if we would have used an infinite number of sub-aperture views, we would get an average of infinitesimal shifts, which looks like a blur to the human eye; yet, sparse light field grid is not enough to produce a smooth, continuous refocused image and therefore the objects outside of the focus plane contain artifacts.

Therefore, in order to produce a continuous refocused image when given a sparse light field data, we had to find a way to add more angular information to the process without actually creating these angles. For this purpose, we employ convolutional neural networks, which have both a strong expressive power on images and generalization abilities. The input to our CNN is a tensor of 2X2 angle views, shifted in advance according to the views' location on the grid and the randomly chosen focus factor. The ground truth is a refocused RGB image generated from a 17X17 light field grid using shifting and averaging. The output of the network is a refocused RGB image (see Fig. 4).

**Ground Truth and Input Generation.** The ground truth images that are used for our network training are created by shifting and averaging the sub-aperture views. We use the Lytro Illum camera datasets in [10, 11], given in the form of lenslet images that can be separated to 14x14 2D angle views. However, we only use 9x9 out of 14x14 given views, to avoid the corruption appearing near the borders of the sub-aperture views located at the edges of the grid, due to vignetting effect. In general, vignetting effects can be removed in pre-processing (as was done with our prototype camera images) and are not the focus of this work.

Still, we notice minor artifacts and bold edges in the resulting image, because the light

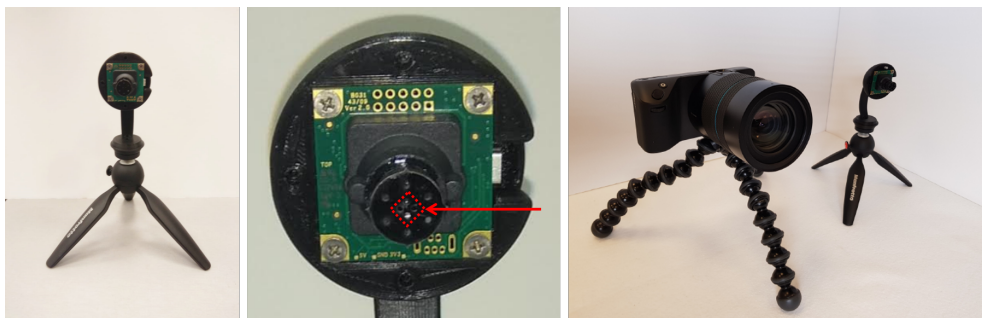


Figure 5: A light field camera prototype built in our lab. Left: the entire setup. Middle: close-up, where the 4 pin-holes are marked in red. Right: our prototype has compact dimensions in comparison to Lytro camera.

field is not dense enough. Therefore, we decide to up-sample the light field grid by adding more samples in-between the given views, while maintaining the original disparity range of the  $9 \times 9$  light field. We generate the in-between views using the algorithm in [14]. We train their network to get an input of  $2 \times 2$  and output a  $3 \times 3$  grid. Then we run the algorithm on  $2 \times 2$  segments of the whole  $9 \times 9$  field, creating a  $17 \times 17$  field (Fig. 1 in sup. material).

In addition, in order to make the artificial refocused ground truth appear as similar as possible to an optically generated refocused image, we only use part of the  $17 \times 17$  input views, which are arranged in a circular shape. Thus, appearing as close as can be to the shape of the camera aperture. This assumption turns out to be correct because the views on the edges of the rectangular grid are often corrupted with black pixels that appear in the process of creating them from the raw light field data. In total, we used 241 views out of the 289 views (see Fig. 2 in sup. material).

The four ( $2 \times 2$ ) sub-aperture images, used as the network’s input, are in the form of a rhombus. We choose this shape, and not a rectangle [14], as it matches a prototype of a light field camera developed in our lab, described next. The chosen views are close enough to the center of the grid but still far enough to take advantage of the disparity between them (see Fig. 2 in sup. material).

**Demo light field camera.** As mentioned earlier, in Lytro camera [28] a single sensor captures both the spatial and angular information. As a result of the inherent tradeoff between them, the sub-aperture views have a low spatial resolution -  $376 \times 541$ . In order to improve the spatial resolution, a light field camera prototype was built in our lab using off-the-shelf components. This prototype favors the spatial resolution ( $1024 \times 1280$ ) over the angular resolution ( $2 \times 2$ ), and a research effort is invested in solving the challenges that arise from the low angular resolution. The camera pin-holes are designed in the shape of a rhombus and a rotating cover exposes a single pin-hole each time. The final setup, which captures 4 sub-aperture images at once, is based on deep reconstruction of compressed light fields [26]. To show the generalization ability of our algorithm across optical devices, we tested it both on images captured by a Lytro camera as well as images captured by our device.

**Network Architecture.** Our proposed refocusing network, **RefocusNet**, is presented in Fig 6. It is a fully convolutional architecture with residual connections, and has two parallel trajectories. Each layer’s output from the first trajectory is the input to the matching layer in the second trajectory. The summation of the seven residuals goes under two operations:

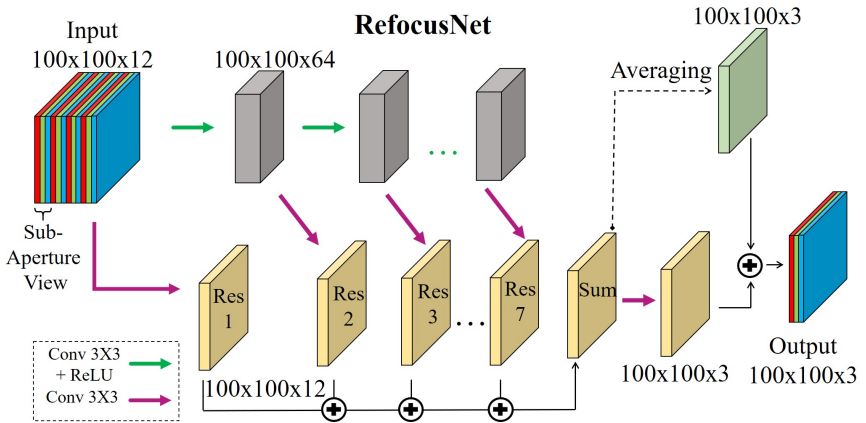


Figure 6: Illustration of the **RefocusNet** architecture. The input consists of 4 sub-aperture views concatenated as a tensor. The network has 2 parallel trajectories; the output of each layer in the first trajectory is used as an input to a matching layer in the second trajectory.

averaging of the 12 channels; and a 3x3 convolution. These two outputs are then summed to get the network output. We train it using the  $\ell_1$  loss as it is more robust to outliers in the training data compared to  $\ell_2$ . The contribution of this specific architecture to our problem is that it preserves the parts in the refocused image that we wish to preserve and repairs the parts that need smoothing e.g. the parts with the ghosting artifacts. While this network is based on an existing architecture named DenoiseNet [81, 82], we have made several important changes in it to improve the performance on both clean and noisy data as shown in the sup. material in several ablation studies that we have conducted.

**Computational efficiency.** In the sup. material we show that an important advantage of our method is that it has low memory usage, and thus can be used in embedded systems.

## 5 Experiments

We turn to compare our technique to other methods: FDL representation based refocusing [14], and two light field synthesis methods [14, 42], which are used to generate a full field of 17x17, followed by the basic shifting and averaging refocusing algorithm [20, 28, 37]. We show our results on the Lytro dataset [1] combined with the dataset of [14], and on our demo camera images. Several ablation studies conducted on our method appear in the sup. material.

### 5.1 Training setup

We train the network on images captured with a Lytro camera, using the Stanford Lytro dataset [1] and the dataset provided in [14]. We divide the data into a training set of 300 images, a validation set of 20 images and a test set of 20 images. The original size of the light field images from both datasets is 376x541x3. However, as described above, we use a synthesizing algorithm [14] to generate a denser light field of 17x17 angle views, which

cropped the pixels on the sides of the images to prevent synthesizing artifacts. Therefore we get smaller input images of size 332x497x3.

We train the network on randomly chosen patches of size 100x100, for  $\sim 300$  epochs with a batch-size of 32 and a total of 9,600 patches per each epoch that are randomly selected from the images. The color channels of each of the input angles are concatenated, i.e., the input tensor in the 4-input-views case is of size [32, 100, 100, 12]. We use the Adam optimizer [15] and the weights are initialized with Xavier initialization [8]. We set the learning rate to be with an exponential decay in a staircase form, where the initial learning rate is set to 0.0005. During inference, we exploit the fact that our network is fully convolutional and therefore can be applied to the whole image at once, although it was trained on patches.

Another issue that emerged during test time is border effects. When using the traditional method [20, 28, 57] of shifting and averaging the sub-aperture views, a linear interpolation is involved in the process to fill in the missing pixels values. However, since at the edges of the image there are no pixel neighbors and the information is missing, the interpolation performed causes border effects. Empirically we have found that there are border effects within 6 pixels from the edges and therefore we remove these parts.

**Focus parameter setting.** For the refocusing done based on the traditional method [20, 28, 57], we used the code in [9], that uses the parameter *pixels* to set the focus point. The relation between the parameter  $\alpha$  from Eq. (4) and *pixels* is:

$$\alpha = \frac{1}{1 - \text{pixels}} \quad (5)$$

The values of the parameter *pixels* were chosen empirically in the range [-1.50,1.30] with spaces of 0.05. These values were selected to match the existing range in the Lytro dataset. Clearly, this range is not suitable for images taken with our demo camera; more details on the latter are provided hereafter.

**Refocusing Noisy Data.** We extend the method described above to noisy light field data. The goal is to generate denoised refocused images when we are given only noisy data. To this end, we have added Gaussian noise with a standard deviation distributed uniformly at random from the range [0, 0.08] to the training data, while the ground truth image is left noiseless. The test is performed on data with the maximal added noise - Gaussian noise with a standard deviation of 0.08.

**Refocusing light field images captured with a demo camera.** As mentioned earlier, to show that our algorithm is independent of the input data type, we tested it also on light field images captured with a prototype designed at our lab. We transferred the network trained for the Lytro dataset to the dataset acquired in our lab, to show the generalization ability of the network. However, for making the transfer possible, the images should have the same scale and therefore we rescaled the demo images by a factor of 0.4 before feeding them to the network; also, the focus range was changed to a new range. The focus parameter *pixels* was determined empirically within the range [-6,6] with spaces of 0.05.

## 5.2 Results

**Lytro data.** Table 1 presents the numerical evaluation of our averaged results on the test set. Our suggested RefocusNet has the best results both in PSNR and SSIM compared to FDL representation based refocusing [17] and the synthesis methods [24] and [22]. The computation time is about 2 orders of magnitude faster than FDL refocusing [17] and 5 orders of magnitude faster than the synthesis methods. Since Wu *et al.* [22] claim that their





Figure 7: Comparison of refocusing and denoising results (PSNR/SSIM) between the traditional method of shifting and averaging with 4 given sub-aperture views ([20], [28], [37]), Kalantari *et al.* [24] and our RefocusNet. Our algorithm cleans the noise in addition to refocusing on the different focus planes, while the traditional method fails in denoising, and the results of Kalantari *et al.* contain some noise (see the car in the blue patch of focus 2.22).

Methods	Input views: 2x2			Input views: 3x3		
	PSNR	SSIM	Time (sec)	PSNR	SSIM	Time (sec)
RefocusNet	<b>46.15</b>	<b>0.997</b>	0.023	<b>48.60</b>	<b>0.998</b>	0.028
FDL [17]	34.10	0.938	$\sim 4$	34.57	0.939	$\sim 4$
Kalantari <i>et al.</i> [24]	43.50	0.990	3,198	44.46	0.993	2,670
Wu <i>et al.</i> [42]	34.49	0.939	1,556	36.35	0.946	1,536

Table 1: Results of light field refocusing with 2x2 and 3x3 given input views using suggested RefocusNet, compared to FDL refocusing [17] and two synthesis methods: Kalantari *et al.* [24] and Wu *et al.* [42]. Since we run FDL refocusing [17] on a CPU, we used the computational times reported in their paper for time estimation on a GPU.

Noisy Data			
	PSNR	SSIM	Time (Sec)
RefocusNet	<b>40.95</b>	<b>0.990</b>	0.022
FDL [17]	29.66	0.777	$\sim 4$
Kalantari <i>et al.</i> [24]	39.45	0.974	3,763

Table 2: Results of light field refocusing on noisy data with 2x2 input views using suggested RefocusNet, compared to FDL refocusing [17] and a synthesis method: Kalantari *et al.* [24].

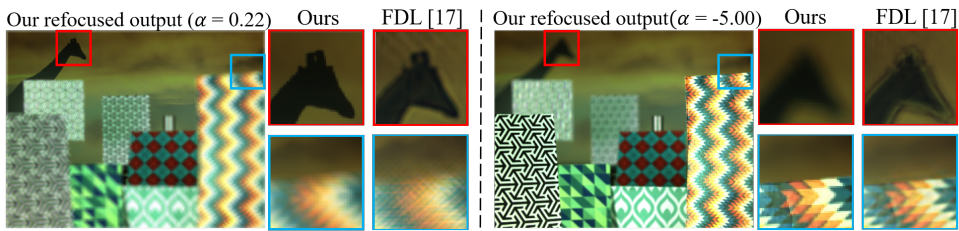


Figure 8: Refocusing results on an image taken with a camera demo built in our lab, comparing our method with FDL refocusing [17]. Note the noticeable artifacts in the background and near the giraffe’s head in their result, while our method successfully smooths these areas.

synthesis algorithm does not perform well when given an angular resolution lower than  $3 \times 3$ , we also trained RefocusNet on an input of  $3 \times 3$  views and compared the results to the other methods. The results show that the performance of all methods improves when increasing the angular resolution, as expected, and our method still has the best performance.

Next, we test our RefocusNet on the noisy test set (described above). Table 2 summarizes the results and demonstrates that RefocusNet has the best results, both in PSNR and SSIM also in this case. Since the traditional refocusing method is based on data averaging, which also has a denoising effect [63], we had to check whether just averaging the 4 given input views achieves satisfying results in the denoising and refocusing tasks. Fig. 7 shows the performance of RefocusNet on a noisy test image, compared to the result of 4 views averaging, Kalantari *et al.* [14] and the clean ground truth. Observe that the shifting and averaging alone leaves noticeable noise, while Kalantari *et al.* removes most of the noise but not all of it, and its computational time is not applicable for real-time applications. However, our algorithm performs better denoising while also refocusing the image, in real-time.

**Real system setup.** Finally, we test our refocusing network on light field images that were taken with a camera demo built in our lab. Notice that no re-training is required and we use the pre-trained network of the Lytro dataset. This demonstrates the generalization ability of our proposed approach. Fig. 8 exhibits our refocusing results compared to FDL refocusing [17], which contain noticeable artifacts. The successful refocusing demonstrated in this example (and others in the sup. material) shows the generalization ability of our technique to different capturing devices.

## 6 Conclusion

In this work, we have proposed an algorithm for digital light field refocusing based on sparse angular information with convolutional networks. To the best of our knowledge, this is the first suggested method for end-to-end light field refocusing based on deep learning. The CNN gets four sub-aperture views and creates a refocused image according to a given focus. We have proposed a novel refocusing network that exhibits very good results in real-time. Experiments show that our method significantly improves the quality of refocusing compared to the previous approaches, with a fast implementation and low memory. Moreover, our approach generalizes well across different sensors without the need to fine-tune it.

## References

- [1] Stanford lytro light field archive. <http://lightfields.stanford.edu/>, .
- [2] Mit computational camera / lightfields. <https://github.com/MITComputationalCamera/LightFields>, .
- [3] Martin Alain, Weston Aenchbacher, and Aljosa Smolic. Interactive light field tilt-shift refocus with generalized shift-and-sum. *arXiv preprint arXiv:1910.04699*, 2019.
- [4] Jonathan T Barron, Andrew Adams, YiChang Shih, and Carlos Hernández. Fast bilateral-space stereo for synthetic defocus. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4466–4474, 2015.
- [5] Ronald N Bracewell. Strip integration in radio astronomy. *Australian Journal of Physics*, 9(2):198–217, 1956.
- [6] Jie Chen and Lap-Pui Chau. Light field compressed sensing over a disparity-aware dictionary. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(4): 855–865, 2017.
- [7] Wenxing Fu, Xi Tong, Chuan Shan, Supeng Zhu, and Baijia Chen. Implementing light field image refocusing algorithm. In *Opto-Electronics and Applied Optics (IEM OPTRONIX), 2015 2nd International Conference on*, pages 1–8. IEEE, 2015.
- [8] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feed-forward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010.
- [9] Mayank Gupta, Arjun Jauhari, Kuldeep Kulkarni, Suren Jayasuriya, Alyosha C Molnar, and Pavan K Turaga. Compressive light field reconstructions using deep learning. In *CVPR Workshops*, pages 1277–1286, 2017.
- [10] Stefan Heber and Thomas Pock. Convolutional networks for shape from light field. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3746–3754, 2016.
- [11] Andre Ivan, In Kyu Park, et al. Synthesizing a 4d spatio-angular consistent light field from a single image. *arXiv preprint arXiv:1903.12364*, 2019.
- [12] Hae-Gon Jeon, Jaesik Park, Gyeongmin Choe, Jinsun Park, Yunsu Bok, Yu-Wing Tai, and In So Kweon. Accurate depth map estimation from a lenslet light field camera. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1547–1555, 2015.
- [13] Hae-Gon Jeon, Jaesik Park, Gyeongmin Choe, Jinsun Park, Yunsu Bok, Yu Wing Tai, and In So Kweon. Depth from a light field image with learning-based matching costs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [14] Nima Khademi Kalantari, Ting-Chun Wang, and Ravi Ramamoorthi. Learning-based view synthesis for light field cameras. *ACM Transactions on Graphics (TOG)*, 35(6): 193, 2016.

- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] Georg Krispel, Michael Opitz, Georg Waltner, Horst Possegger, and Horst Bischof. Fuseseg: Lidar point cloud segmentation fusing multi-modal data. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 1874–1883, 2020.
- [17] Mikael Le Pendu, Christine Guillemot, and Aljosa Smolic. A fourier disparity layer representation for light fields. *IEEE Transactions on Image Processing*, 28(11):5740–5753, 2019.
- [18] Ensun Lee, Seohee Yang, Miseon Han, and Jeongtae Kim. Depth-based refocusing for reducing directional aliasing artifacts. *Optics express*, 24(24):28065–28079, 2016.
- [19] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42. ACM, 1996.
- [20] Marc Levoy, Billy Chen, Vaibhav Vaish, Mark Horowitz, Ian McDowall, and Mark Bolas. Synthetic aperture confocal imaging. *ACM Transactions on Graphics (TOG)*, 23(3):825–834, 2004.
- [21] Chia-Kai Liang, Tai-Hsu Lin, Bing-Yi Wong, Chi Liu, and Homer H Chen. Programmable aperture photography: multiplexed light field acquisition. *ACM Transactions on Graphics (TOG)*, 27(3):55, 2008.
- [22] José G Marichal-Hernández, Jonas P Lüke, Fernando L Rosa, and José M Rodríguez-Ramos. Fast approximate 4d: 3d discrete radon transform, from light field to focal stack with  $\mathcal{O}(n^4)$  sums. In *Real-Time Image and Video Processing 2011*, volume 7871, page 78710G. International Society for Optics and Photonics, 2011.
- [23] Kshitij Marwah, Gordon Wetzstein, Yosuke Bando, and Ramesh Raskar. Compressive light field photography using overcomplete dictionaries and optimized projections. *ACM Transactions on Graphics (TOG)*, 32(4):46, 2013.
- [24] Kevin Matzen, Vivien Nguyen, Dillon Yao, You Zhang, Ren Ng, et al. Synthetic defocus and look-ahead autofocus for casual videography. *arXiv preprint arXiv:1905.06326*, 2019.
- [25] Kaushik Mitra and Ashok Veeraraghavan. Light field denoising, light field superresolution and stereo camera based refocussing using a gmm light field patch prior. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2012 IEEE Computer Society Conference on*, pages 22–28. IEEE, 2012.
- [26] Ofir Nabati, David Mendlovic, and Raja Giryes. Fast and accurate reconstruction of compressed color light field. In *Computational Photography (ICCP), 2018 IEEE International Conference on*, pages 1–11. IEEE, 2018.
- [27] Ren Ng. Fourier slice photography. *ACM transactions on graphics (TOG)*, 24(3):735–744, 2005.

- [28] Ren Ng, Marc Levoy, Mathieu Brédif, Gene Duval, Mark Horowitz, Pat Hanrahan, et al. Light field photography with a hand-held plenoptic camera. *Computer Science Technical Report CSTR*, 2(11):1–11, 2005.
- [29] Jinsun Park, Yu-Wing Tai, Donghyeon Cho, and In So Kweon. A unified approach of multi-scale deep and hand-crafted features for defocus estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1736–1745, 2017.
- [30] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM Transactions on Graphics (TOG)*, 36(6):235, 2017.
- [31] T. Remez, O. Litany, R. Giryes, and A. M. Bronstein. Deep class-aware image denoising. In *IEEE International Conference on Image Processing (ICIP)*, pages 1895–1899, Sep. 2017.
- [32] T. Remez, O. Litany, R. Giryes, and A. M. Bronstein. Class-aware fully convolutional gaussian and poisson denoising. *IEEE Transactions on Image Processing*, 27(11):5707–5722, Nov 2018.
- [33] Inderpreet Singh and Nirvair Neeru. Performance comparison of various image denoising filters under spatial domain. *International Journal of Computer Applications*, 96(19), 2014.
- [34] Pratul P Srinivasan, Tongzhou Wang, Ashwin Sreelal, Ravi Ramamoorthi, and Ren Ng. Learning to synthesize a 4d rgb-d light field from a single image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2243–2251, 2017.
- [35] Yuxiang Sun, Weixun Zuo, and Ming Liu. Rtfnet: Rgb-thermal fusion network for semantic segmentation of urban scenes. *IEEE Robotics and Automation Letters*, 4(3): 2576–2583, 2019.
- [36] Supasorn Suwajanakorn, Carlos Hernandez, and Steven M Seitz. Depth from focus with your mobile phone. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3497–3506, 2015.
- [37] Vaibhav Vaish, Bennett Wilburn, Neel Joshi, and Marc Levoy. Using plane+ parallax for calibrating dense camera arrays. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I. IEEE, 2004.
- [38] Ashok Veeraraghavan, Ramesh Raskar, Amit Agrawal, Ankit Mohan, and Jack Tumblin. Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM transactions on graphics (TOG)*, 26(3):69, 2007.
- [39] Lijun Wang, Xiaohui Shen, Jianming Zhang, Oliver Wang, Zhe Lin, Chih-Yao Hsieh, Sarah Kong, and Huchuan Lu. DeepLens: shallow depth of field from a single image. *arXiv preprint arXiv:1810.08100*, 2018.
- [40] Yunlong Wang, Fei Liu, Zilei Wang, Guangqi Hou, Zhenan Sun, and Tieniu Tan. End-to-end view synthesis for light field imaging with pseudo 4d-cnn. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 333–348, 2018.

- [41] 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. *ACM Transactions on Graphics (TOG)*, 24(3): 765–776, 2005.
- [42] Gaochang Wu, Mandan Zhao, Liangyong Wang, Qionghai Dai, Tianyou Chai, and Yebin Liu. Light field reconstruction using deep convolutional network on epi. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2017, page 2, 2017.
- [43] Zhaolin Xiao, Qing Wang, Guoqing Zhou, and Jingyi Yu. Aliasing detection and reduction in plenoptic imaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3326–3333, 2014.