

# A Study of Efficient Light Field Subsampling and Reconstruction Strategies

Yang Chen, Martin Alain, and Aljosa Smolic

*V-SENSE project*  
*Graphics Vision and Visualisation group (GV2)*  
*Trinity College Dublin*\*

## Abstract

Limited angular resolution is one of the main obstacles for practical applications of light fields. Although numerous approaches have been proposed to enhance angular resolution, view selection strategies have not been well explored in this area. In this paper, we study subsampling and reconstruction strategies for light fields. First, different subsampling strategies are studied with a fixed sampling ratio, such as row-wise sampling, column-wise sampling, or their combinations. Second, several strategies are explored to reconstruct intermediate views from four regularly sampled input views. The influence of the angular density of the input is also evaluated. We evaluate these strategies on both real-world and synthetic datasets, and optimal selection strategies are devised from our results. These can be applied in future light field research such as compression, angular super-resolution, and design of camera systems.

**Keywords:** Light Fields, Angular Super-resolution, Reconstruction and Subsampling, View Synthesis

## 1 Introduction

A light field (LF) is described as all light rays passing through a given 3D volume in the pioneer work of Levoy et al. [Levoy and Hanrahan, 1996]. Different from traditional 2D imaging systems, a 4D LF imaging system captures not only two spatial but also two additional angular dimensions. The two plane representation of light fields is adopted in this paper, which is represented as a collection of views taken from several view points parallel to a common plane, as shown in Figure 1. Another common representation of light fields are Epipolar Plane Images (EPI), which are 2D slices of the 4D light field obtained by fixing one spatial and one angular dimension.

Dense light fields are advantageous for various applications including virtual reality, augmented reality and image-based rendering [Yu, 2017, Wu et al., 2017a]. However, dense light fields are usually difficult to capture in practice and a trade-off has to be found between spatial and angular resolution. For example, camera arrays usually have good spatial resolution but sparse angular sampling, while the opposite is usually the case for plenoptic cameras.




Figure 1: Subsampling and reconstruction of a two-plane representation light field.

\*This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under the Grant Number 15/RP/2776. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

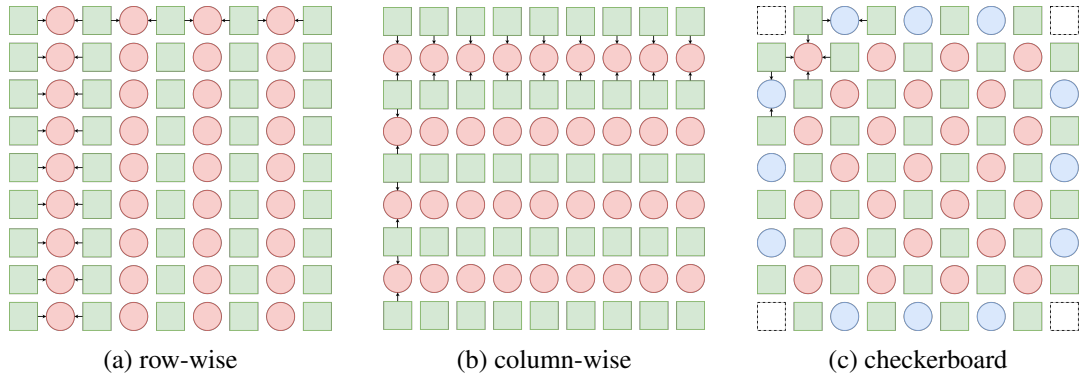


Figure 2: Three basic subsampling strategies. The squares are sampled views, the circles are reconstructed views and the dashed squares are unused views. Blue circles in (c) are reconstructed by row-wise or column-wise from two adjacent views to complete the LF.

Spatial super-resolution of light fields is a widely studied problem and many methods have been proposed with impressive results [Alain and Smolic, 2018, Rossi and Frossard, 2017, Yoon et al., 2015]. As for the angular resolution enhancement in light fields, one common solution is to apply view synthesis methods on LFs [Yoon et al., 2015, Kalantari et al., 2016, Wang et al., 2018, Vagharshakyan et al., 2017]. However, only limited attention has been paid to view selection strategies and previous work usually arbitrarily selects one fixed input pattern, such as along one angular row or within a  $N \times N$  square matrix of views.

In this paper, we investigate light field subsampling and reconstruction strategies and evaluate their performance. We identify three issues to face: First, a benchmark method is required to compare these strategies. State-of-the-art (SOTA) light field subsampling and reconstruction methods are reviewed and evaluated experimentally, and the best performing approach is chosen as the benchmark for following experiments. Second, since LFs contain huge amounts of data, which poses challenges for storage and transmission, efficient subsampling methods are needed for LFs. Given a dense LF as input, we investigate which subsampling strategy can produce the best reconstruction results while keeping a satisfying sampling ratio. Third, since dense LFs are expensive to capture, reconstructing a dense LF from sparse input becomes an important and challenging topic. Here we investigate different strategies and different levels of sparsity for the task of reconstructing a LF from corner views. Eventually, we derive optimal strategies for each problem based on our results.

This paper is organized as follows. In Section 2, we review existing methods for light field angular super-resolution and more general techniques for view synthesis and video frame interpolation. In Section 3, a benchmark method is selected and the experiments for evaluation of subsampling and reconstruction strategies are described in detail. Then, these designs are evaluated and compared in Section 4 after applying the selected benchmark method. Finally, we present our conclusions in Section 5.

## 2 Related Work

**Angular Super-resolution for Light Fields** Wanner and Goldluecke proposed a variational framework to generate novel views from sparse input views, of which ghost artifacts can still be observed on the synthesized results [Wanner and Goldluecke, 2013]. Densely sampled light fields have small disparity between adjacent views, which makes them very suitable for frequency domain analysis. Therefore, Shi et al. proposed to reconstruct the dense light field by optimizing frequency coefficients in the continuous 2D Fourier domain [Shi et al., 2014]. A more advanced shearlet transform is adapted in [Vagharshakyan et al., 2017] on Epipolar Plane Images. However, wide EPI input is required by these Fourier methods, which is not feasible for some applications.

The strength of deep learning based approaches has been demonstrated in several LF super-resolution methods. A first example designed a multi-stream network for light field spatial and angular super-resolution [Yoon et al., 2015]. Kalantari et al. proposed a learning based framework for light field view synthesis, which requires depth estimation as an intermediate step [Kalantari et al., 2016]. A sparse input consisting of the cor-

Table 1: Comparison between SOTA LF interpolation methods

row-wise	PSNR/SSIM		
	Mean	HCI	Stanford
Linear	32.37/0.9464	31.82/0.9501	33.47/0.9390
Shearlet [Vagharshakyan et al., 2017]	34.92/0.9592	35.78/0.9750	33.21/0.9277
LFEPi [Wu et al., 2017b]	37.39/0.9451	37.51/0.9420	37.15/0.9515
SepConv [Niklaus et al., 2017]	<b>38.94/0.9910</b>	<b>39.38/0.9945</b>	<b>38.07/0.9839</b>

Table 2: Comparison of pretrained, fine-tuned and retrained SepConv for row-wise interpolation

row-wise	pretrained	Fine-tuned	Retrained
PSNR(dB)	38.08	<b>39.74</b>	39.41
SSIM	0.9893	<b>0.9925</b>	0.9923

ner views of a dense light field is fed into a convolutional neural network to synthesize the original intermediate views. Wu et al. re-modeled light field angular super-resolution as a detail restoration problem in 2D EPI space (LFEPi) [Wu et al., 2017b]. A blur-restoration-deblur framework is built to process EPIs of a sparsely sampled LF and to recover the angular detail with a convolutional neural network. To utilize inherent consistency of the LF, Wang et al. introduced a pseudo 4D convolution by combining a 2D convolution on EPIs and a sequential 3D convolution [Wang et al., 2018]. While these methods presented innovative ideas and good results, fixed subsampling and reconstruction strategies are applied without investigating and optimising alternatives.

**View Synthesis and Video Frame Interpolation** Existing depth image-based rendering and video frame interpolation methods can be directly applied to the LF angular super-resolution problem. A deep architecture was introduced by Flynn et al. to synthesize novel views from wide disparity real-world inputs [Flynn et al., 2016]. Niklaus et al. proposed pairs of spatially individual 1D kernels, which are estimated from a trained convolutional neural network, to estimate motion and color interpolation within one stage (SepConv) [Niklaus et al., 2017].

### 3 Study of efficient subsampling and reconstruction strategies

In this section, we investigate different strategies for light field subsampling and reconstruction. Firstly, a benchmark LF view interpolation method has to be selected from SOTA approaches to evaluate all strategies. Next, given a fixed sampling ratio, three light field subsampling strategies are studied to reconstruct full size LFs from each sampled LF (Figure 2). Finally, six different reconstruction strategies are explored to generate a dense light field from inputs of varying sparsity (Figures 3 & 4).

#### 3.1 Benchmark method selection

As our goal is to investigate sampling and reconstruction strategies for LFs, we first select a benchmark method to be applied in our experiments. The benchmark has to be flexible to work in different configurations but should also provide best possible interpolation quality. We therefore evaluate SepConv [Niklaus et al., 2017], Shearlet [Vagharshakyan et al., 2017] and LFEPi [Wu et al., 2017b] in an initial study. LFEPi is a representative learning based LF view synthesis method, and Shearlet is an efficient non-learning based reconstruction method in the Fourier domain. SepConv [Niklaus et al., 2017] was initially designed for video frame interpolation and employs a neural network based kernel estimator to interpolate views between adjacent input views. As such it is very flexible and can be also be used in various ways of LF view interpolation.

These SOTA methods are evaluated by using the same input pattern shown in Figure 2a resulting from row-wise sampling, in which sampled input views are represented as green squares and reconstructed output views are represented as red circles. According to Table 1, SepConv numerically outperforms all other methods significantly. Shearlet achieves better performance than linear interpolation. LFEPi scores well on PSNR, while Shearlet performs better in terms of average SSIM. However, both these two methods require an EPI as input. Thus SepConv is not only the best performing approach, but also the only one that can easily be adapted to

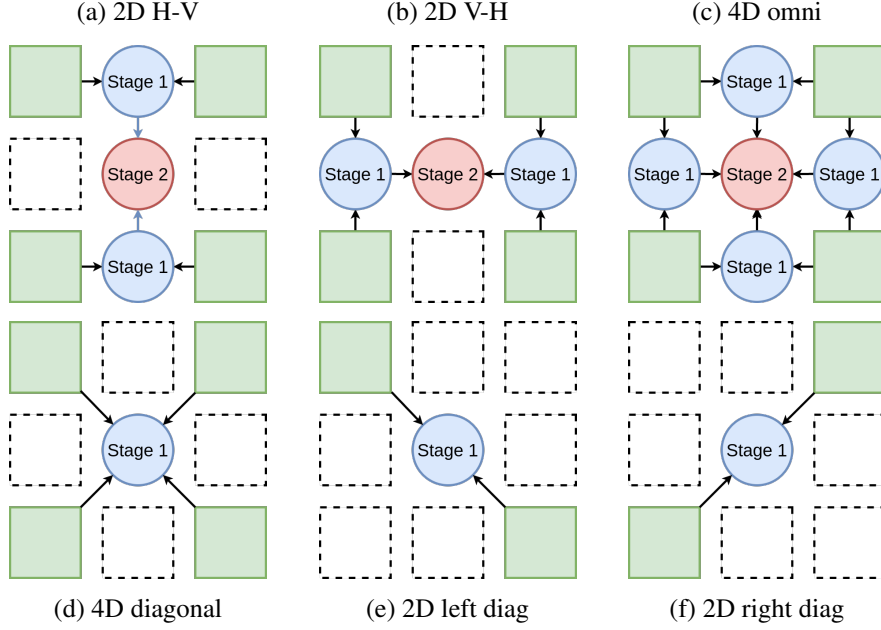


Figure 3: Six reconstruction strategies to interpolate 3x3 views from 4 input corner views (green), dashed square views are not used for interpolation, different colors of circles identify output views from different stages

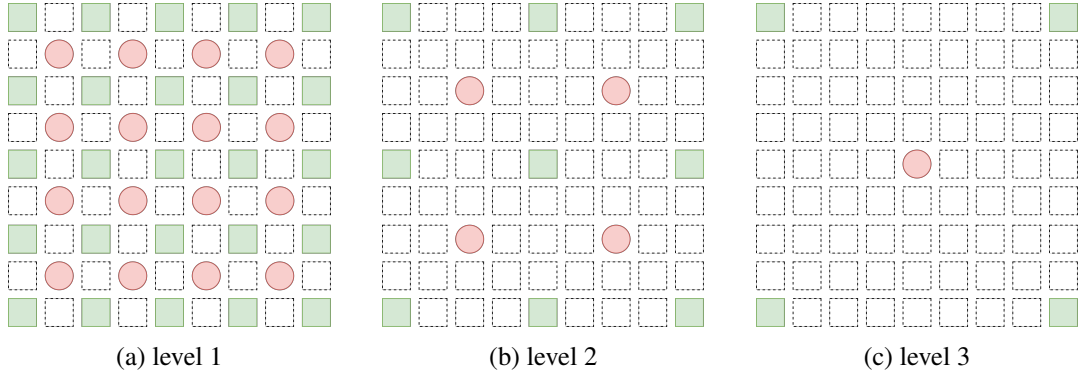


Figure 4: Three levels of angular density for LF reconstruction. The *IRs* of these levels are 30.9%, 11.1% and 4.9% respectively.

different configurations, as it only needs a pair of RGB images as input. We therefore continue to use SepConv in our further experiments.

Since SepConv was originally trained for video frame interpolation, we additionally fine-tuned the pre-trained model on our light field training dataset in order to further improve the performance. Table 2 shows improvements we can achieve by fine-tuning and retraining the initial network. For some of our experiments detailed below, we have to retrain SepConv appropriately in order to work with more than 2 input views, e.g. left, right, top and bottom neighbors.

### 3.2 Study of basic light field subsampling strategies

In this set of experiments, we compare basic subsampling strategies as illustrated in Figure 2, with the goal to identify the most efficient basic subsampling strategy among row-wise, column-wise and checkerboard. To evaluate these strategies, the sampled LFs are reconstructed using view interpolation. After reconstruction, the quality of the synthesized views can be compared to the corresponding ground truth. All these strategies have the same ratio of sampled views to total views, which is an important measure of the sparsity and defined as the *InputRatio* (*IR*) in equation (1):

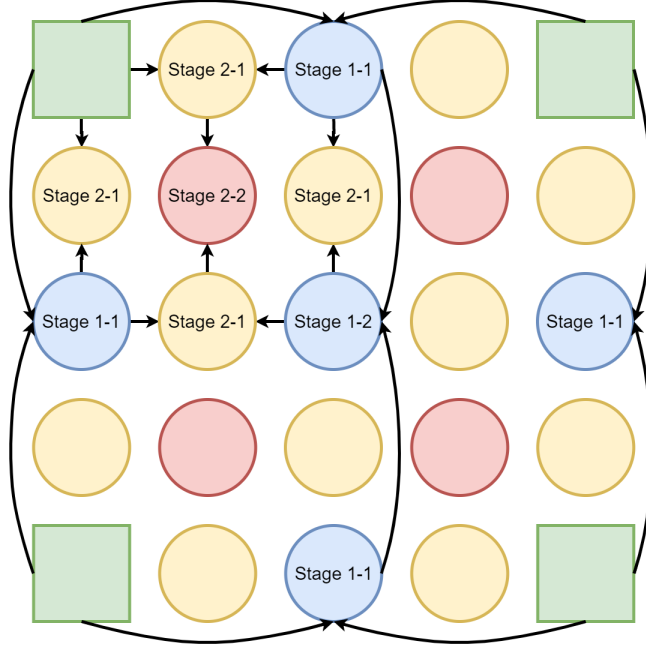


Figure 5: Stages to complete a quarter of the full LF for an input angular density of level 2 with  $IR = 11.1\%$  (see Fig. 4)

$$IR = \frac{N_{InputViews}}{N_{InputViews} + N_{OutputViews}} \quad (1)$$

where the numbers of input views and output views of the interpolation method are  $N_{InputViews}$  and  $N_{OutputViews}$ , respectively. The total number of views of the completed light field can be represented as the sum of  $N_{InputViews}$  and  $N_{OutputViews}$ . The three basic subsampling strategies as shown in Figure 2 all have  $IR \approx 55\%$ .

We applied our fine-tuned SepConv as explained in Section 3.1 to interpolate the necessary views for row-wise and column-wise strategies from neighbouring views. For the checkerboard pattern we had to modify the original SepConv network in order to accept 4 views as input, top, bottom, left and right. The outer views depicted as blue circles in Figure 2c were synthesized by row-wise or column-wise interpolation from 2 neighbouring views.

### 3.3 Study of sparse light field reconstruction strategies

In this set of experiments, we compare different reconstruction strategies for sparse light fields. Taking a  $3 \times 3$  matrix of views as example, six progressive reconstruction strategies can be applied as presented in Figure 3. Using four corner images as input, we can reconstruct the side images using the same row-wise and column-wise interpolation as before. Thus, the question becomes which is the best way to reconstruct the central view. Three of these strategies, including 2D horizontal-vertical (2D H-V), 2D vertical-horizontal (2D V-H) and 4D omni, are two stage cases involving generating side views as intermediate stage. The other three, including 4D diagonal, 2D left diagonal and 2D right diagonal, require only one stage to generate the central view.

To further study the influence of the angular density of the input views, we investigate three levels of density as shown in Figure 4. The distance between two input views is 1 view in level 1, and there are 3 views and 7 views distance in level 2 and 3, respectively. To compare fully reconstructed light fields, we reconstruct all missing views using appropriate strategies as follows. Level 1 is completed by row-wise and column-wise interpolation of the side views, using the 2D H-V method unless otherwise specified. The choice of 2D H-V as the default method is justified by its better performance demonstrated in Section 4. For level 2 and 3, we recursively fill using lower level methods to complete the full scale LF. The reconstruction of a quarter of the

Table 3: Evaluation of three basic subsampling strategies from Figure 2

	row-wise		column-wise		checkerboard	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
HCI	40.49	0.9956	40.33	0.9958	39.24	0.9935
Lytro	39.32	0.9932	39.07	0.9925	38.67	0.9921
Stanford	40.22	0.9936	39.37	0.9924	39.47	0.9928
Mean	<b>39.74</b>	<b>0.9925</b>	39.20	0.9915	39.33	0.9918

Table 4: Evaluation of six reconstruction strategies from Figure 3

	PSNR	SSIM
2D H-V	<b>37.42</b>	<b>0.9884</b>
2D V-H	37.21	0.9881
2D left diagonal	35.77	0.9846
2D right diagonal	35.86	0.9842
4D omni	36.51	0.9863
4D diagonal	36.86	0.9838

LF is shown in Figure 5, which is applied iteratively to each quarter one by one. Please find more detailed completion stages in the supplementary material <sup>1</sup>.

## 4 Results

In this section, we summarize the results of our experiments, which were performed on an Intel Core i7-6700k 4.0GHz CPU, while neural network refining was performed on Nvidia Titan Xp GPUs.

The performance of strategies is evaluated on both real-world and synthetic LF datasets to validate their robustness. We used 27 real-world light fields captured by Lytro Illum cameras provided by EPFL [Rerabek and Ebrahimi, 2016] and INRIA [INR, ], and 11 light fields from the Stanford dataset taken by a camera gantry [Sta, ]. As for the synthetic LF dataset, all 28 light fields from the HCI benchmark [Honauer et al., 2016] were used. 10 light fields in total were selected from these datasets as the test set and the rest as training set. Additionally, 160 light fields from LF intrinsic [Alperovich et al., 2018] were added for retraining of SepConv to avoid the overfitting. All views were cropped to equal 512x512 resolution to accelerate the computation. In this study, we use the peak signal-to-noise ratio (PSNR) and the structural similarity (SSIM) on RGB images to evaluate the algorithms numerically. For each light field, unless emphasized specifically, mean numerical results are computed over all views of the full light fields.

The results of the evaluation of basic subsampling strategies are shown in Table 3. Row-wise interpolation achieves the best scores on most light fields over all datasets. Please view our website for comprehensive results of each light field <sup>1</sup>. The difference to column-wise is most prominent for the Stanford data which was captured by a gantry, resulting in unequal vertical displacement. For the checkerboard pattern, the retraining of SepConv to accept 4 views may be the reason of its worse performance, as this way it could not benefit from the pretrained model of SepConv.

The results of sparse LF reconstruction strategies are shown in Table 4. Again, 4D strategies use a retrained model while the others use a fine-tuned model. The central view is used to evaluate these results as it is the only one that gets always reconstructed using any of the six strategies when filling a 3x3 block of views. 2D H-V performs best, as it accumulates less error than other strategies with 2 stages (row-wise first gives best reference for stage 2), while it has smaller distance between input views compared to direct diagonal strategies. Visual results of reconstruction strategies are shown in Figure 6. Occlusion artifacts around the tip of sword can be observed in diagonal strategies.

Finally, the effect of different levels of angular density is studied in Table 5. Since 2D H-V is the optimal strategy according to the previous conclusion, it is utilized to complete the full scale LFs recursively as explained before. All synthesized views are averaged in this evaluation (different from only central views in Table 4). From Table 5 we get the expected decrease of interpolation quality with sparsity.

These insights about sparsity vs. quality and the best subsampling and reconstruction strategies can be beneficial for the design of LF coding approaches (maximum quality that can be achieved when omitting views) or camera systems (maximum camera distance for a desired quality and density).

<sup>1</sup><https://v-sense.scss.tcd.ie/?p=4450>

Table 5: Comparison between level 1 & 2 & 3 from Figure 4 using 2D H-V

2D H-V	Level 1	Level 2	Level 3
PSNR(dB)	37.34	35.12	31.79
SSIM	0.9886	0.9825	0.9613




Figure 6: Visual results of six reconstruction strategies from Figure 3

## 5 Conclusion

In this paper, we presented a comprehensive study comparing different strategies for efficient light field subsampling and reconstruction. For this purpose we selected an existing view synthesis method among the best performing state-of-the-art techniques. Using this benchmark method, we first evaluate the best subsampling approach among a row-wise, a column-wise, and a checkerboard pattern with a fixed interpolation ratio, which concludes that the row-wise approach offers the best performances. Second, we investigate corner-based central view generation and compare the performance of six possible reconstruction strategies. In addition, we evaluate a multi-stage approach to reconstruct dense light fields from subsampled input with different levels of angular density. We found that using a row-wise followed by a col-wise reconstruction yields best performance. We hope these findings will help inspire researches related to light field subsampling and reconstruction, such as compression and camera array design. Further, a more explicit analysis regarding the relation of the disparity range to the strategy selection could be performed.

## References

- [INR, ] Inria Lytro Illum dataset. <http://www.irisa.fr/temics/demos/lightField/CLIM/DataSoftware.html>. accessed: 26-01-2018.
- [Sta, ] The stanford light field archive. <http://lightfield.stanford.edu/lfs.html>. accessed: 05-03-2019.
- [Alain and Smolic, 2018] Alain, M. and Smolic, A. (2018). Light field super-resolution via LFBM5D sparse coding. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 2501–2505. IEEE.

- [Alperovich et al., 2018] Alperovich, A., Johannsen, O., Strecke, M., and Goldluecke, B. (2018). Light field intrinsics with a deep encoder-decoder network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9145–9154.
- [Flynn et al., 2016] Flynn, J., Neulander, I., Philbin, J., and Snavely, N. (2016). DeepStereo: Learning to predict new views from the world’s imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5515–5524.
- [Honauer et al., 2016] Honauer, K., Johannsen, O., Kondermann, D., and Goldluecke, B. (2016). A dataset and evaluation methodology for depth estimation on 4D light fields. In *Asian Conference on Computer Vision*. Springer.
- [Kalantari et al., 2016] Kalantari, N. K., Wang, T.-C., and Ramamoorthi, R. (2016). Learning-based view synthesis for light field cameras. *ACM Transactions on Graphics (TOG)*, 35(6):1–10.
- [Levoy and Hanrahan, 1996] Levoy, M. and Hanrahan, P. (1996). Light field rendering. In *Proc. SIGGRAPH*, pages 31–42.
- [Niklaus et al., 2017] Niklaus, S., Mai, L., and Liu, F. (2017). Video frame interpolation via adaptive separable convolution. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 261–270.
- [Rerabek and Ebrahimi, 2016] Rerabek, M. and Ebrahimi, T. (2016). New light field image dataset. In *Proceedings of the International Conference on Quality of Multimedia Experience*.
- [Rossi and Frossard, 2017] Rossi, M. and Frossard, P. (2017). Graph-based light field super-resolution. In *2017 IEEE 19th International Workshop on Multimedia Signal Processing (MMSP)*, pages 1–6. IEEE.
- [Shi et al., 2014] Shi, L., Hassanieh, H., Davis, A., Katabi, D., and Durand, F. (2014). Light field reconstruction using sparsity in the continuous fourier domain. *ACM Transactions on Graphics (TOG)*, 34(1):1–13.
- [Vagharshakyan et al., 2017] Vagharshakyan, S., Bregovic, R., and Gotchev, A. (2017). Light field reconstruction using shearlet transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(1):133–147.
- [Wang et al., 2018] Wang, Y., Liu, F., Wang, Z., Hou, G., Sun, Z., and Tan, T. (2018). End-to-end view synthesis for light field imaging with pseudo 4DCNN. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 333–348.
- [Wanner and Goldluecke, 2013] Wanner, S. and Goldluecke, B. (2013). Variational light field analysis for disparity estimation and super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3):606–619.
- [Wu et al., 2017a] Wu, G., Masia, B., Jarabo, A., Zhang, Y., Wang, L., Dai, Q., Chai, T., and Liu, Y. (2017a). Light field image processing: An overview. *IEEE Journal of Selected Topics in Signal Processing*, 11(7):926–954.
- [Wu et al., 2017b] Wu, G., Zhao, M., Wang, L., Dai, Q., Chai, T., and Liu, Y. (2017b). Light field reconstruction using deep convolutional network on EPI. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6319–6327.
- [Yoon et al., 2015] Yoon, Y., Jeon, H.-G., Yoo, D., Lee, J.-Y., and So Kweon, I. (2015). Learning a deep convolutional network for light-field image super-resolution. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 24–32.
- [Yu, 2017] Yu, J. (2017). A light-field journey to virtual reality. *IEEE MultiMedia*, 24(2):104–112.