

Light Field Denoising by Sparse 5D Transform Domain Collaborative Filtering

Martin Alain
V-SENSE project

Graphics Vision and Visualisation group (GV2)
Trinity College Dublin
alainm@scss.tcd.ie

Aljosa Smolic
V-SENSE project

Graphics Vision and Visualisation group (GV2)
Trinity College Dublin
smolica@scss.tcd.ie

Abstract—In this paper, we propose to extend the state-of-the-art BM3D image denoising filter to light fields, and we denote our method LFBM5D. We take full advantage of the 4D nature of light fields by creating disparity compensated 4D patches which are then stacked together with similar 4D patches along a 5th dimension. We then filter these 5D patches in the 5D transform domain, obtained by cascading a 2D spatial transform, a 2D angular transform, and a 1D transform applied along the similarities. Furthermore, we propose to use the shape-adaptive DCT as the 2D angular transform to be robust to occlusions. Results show a significant improvement in synthetic noise removal compared to state-of-the-art methods, for both light fields captured with a lenslet camera or a gantry. Experiments on Lytro Illum camera noise removal also demonstrate a clear improvement of the light field quality.

I. INTRODUCTION

Light fields aim to capture all light rays passing through a given amount of the 3D space [1]. Compared to traditional 2D imaging systems which capture the spatial intensity of light rays, a 4D light field also contains the angular direction of the rays. We adopt in this paper the common two-plane parametrization, and a light field can be formally represented as a 4D function $\Omega \times \Pi \rightarrow \mathbb{R}$, $(x, y, s, t) \rightarrow L(x, y, s, t)$ in which the plane Ω represents the spatial distribution of light rays, indexed by (x, y) , while Π corresponds to their angular distribution, indexed by (s, t) . Perhaps the easiest way to visualize a light field is to consider it as a matrix of views (see Fig. 1), also called sub-aperture images (SAI). Each SAI represents a 2D slice of the light field over the spatial dimensions (x, y) . 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 (sx - or yt -planes, see Fig. 1).

Applications of light field include refocusing of an image after capture [2], rendering of virtual points of view [1], [3], or depth estimation [4]. Many camera architectures have been proposed to capture light fields, including gantry [5], camera arrays [6], or more recently hand-held lenslet cameras [2], [7]. However, light field captured with the latter cameras exhibit visible noise artifacts. Thus light field noise removal has become an important task for aesthetic purposes or as a pre-processing step. Formally, we define the noisy light field as follows:

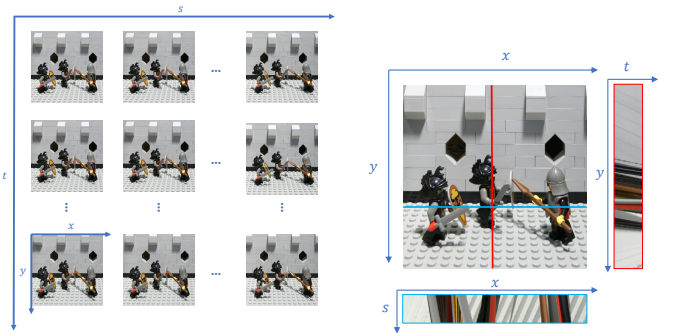


Fig. 1. Examples of light field representations: matrix of sub-aperture images (SAI) (left); and Epipolar Plane Images (EPI) (right) shown below and on the right of the center SAI.

$$L^{noisy} = L + \eta \quad (1)$$

where L is the source light field that we want to reconstruct, L^{noisy} is the observed data, and η corresponds to Additive White Gaussian Noise (AWGN) with variance σ^2 .

Several methods have been proposed in the past years to address light field denoising [8]–[12], which benefit from the very active research conducted on single image denoising [13]–[19]. However, the best performing methods of the state-of-the-art [12] are obtained by processing 2D slices of the light field (SAI or EPI) stacked into a 3rd dimension, and do not fully exploit the 4D structure of light fields.

In this paper, we propose to extend the BM3D filter [16], which exploits natural local redundancies existing in natural images, by taking into account the 2D redundancies occurring over the angular dimensions of the light field. We first propose to create 4D disparity compensated patches, which contain highly redundant 2D patches, and stack them together with similar 4D patches over a 5th dimension. The 5D patches obtained in this fashion are then processed in the 5D transform domain, where the underlying true signal is sparsely represented. The 5D transform is a combination of a 2D spatial transform, a 2D angular transform, and a 1D transform applied on the 5th dimension corresponding to similarities. In order to be robust to occlusions and maintain the sparsity of the 5D spectrum, we propose to use the shape-adaptive DCT (SADCT) as the 2D angular transform. Furthermore, our method does not make assumption on the angular sampling

density, and performs equally on dense or sparsely sampled light fields.

This paper is organized as follows. In section II we review the existing work on light field denoising. Section III describes more in detail the proposed LFBM5D approach. Finally, we evaluate in section IV the denoising performances of our approach against state-of-the-art methods.

II. BACKGROUND

A naive approach to light field denoising consists in applying existing image denoising methods to the SAIs independently. We refer the reader to [18], [19] for a survey of image denoising methods. To better take into account the redundancies existing in-between the different SAIs, one can instead process the EPIs. In [9], a two-step method is proposed which first denoises EPIs taken along a given spatial and angular dimension (e.g. sx -plane), and then processes this first estimate using the complementary EPIs (e.g. yt -plane).

However, these methods only consider 2D slices of the 4D light field. To further exploit the inter-dependencies of the light field, it is also possible to stack the SAIs or EPIs and filter along this additional dimension. Among the existing methods, the best denoising performances are thus obtained by applying the VBM4D video filter [20] on such sequences [12].

Other methods propose to take into account the 4D structure of the light field through a GMM light field patch prior [8], variational regularization [10], or a global 4D transform domain hyperfan filter [11], denoted HF4D in this paper. Although these methods fully exploit the 4D structure of the light field, they fail to outperform the VBM4D filter.

The VBM4D filter is an extension of the BM3D filter [16] to video denoising. We refer the reader to [17] for a more detailed analysis of the BM3D filter. As BM3D is still among the best image denoising methods [18], we propose in this paper to extend it to light fields, and we denote our approach LFBM5D. We show that by fully taking into account the 4D light field structure, we are able to significantly outperform the best state-of-the-art methods.

III. LIGHT FIELD DENOISING IN THE 5D TRANSFORM DOMAIN

In this section, we describe our method called LFBM5D. Note that we adopt notations similar to [17].

A. Overview

Our method follows the same two-step procedure as BM3D: first a so-called basic estimate is obtained after hard-thresholding the transform coefficients of the 5D patches, then the basic estimate is used to perform a Wiener filtering in the 5D transform domain and produce the final denoised estimate. The steps are denoted *hard* and *wiener* respectively.

To process the noisy light field L^{noisy} , we iterate over the sub-aperture images (SAIs). First a specific SAI, denoted as the reference SAI, is selected. At initialization, we select the central SAI of the light field. We then consider the neighboring SAIs found in a so-called angular search window of size $n_a \times$

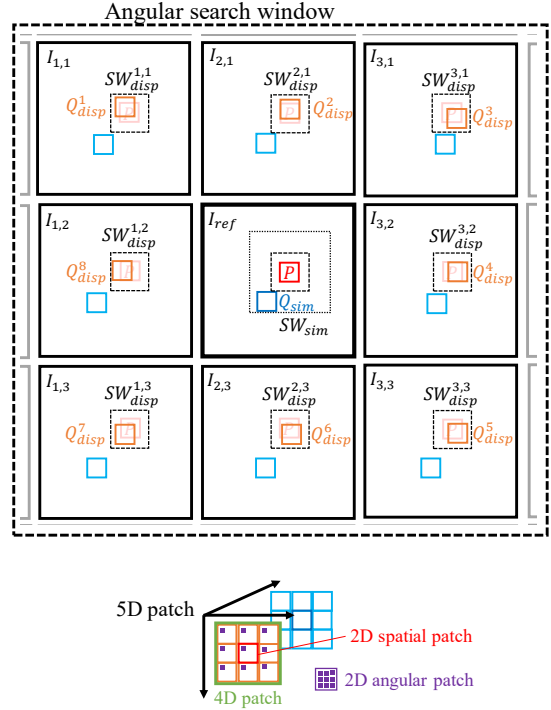


Fig. 2. The SAIs of a light field are iteratively processed by considering a reference SAI and its neighboring SAIs in a so-called angular search window. When processing SAIs in an angular search window, 4D patches are obtained by taking disparity compensated 2D patches in the neighboring SAIs with respect with a 2D patch in the reference SAI. 5D patches are finally built by stacking similar 4D patches along a 5th dimension.

n_a , and filter all these SAIs together. For that purpose, we first iterate over the 2D patches of the reference SAI. As explained in the next section, 5D patches are created by combining 2D patches of the reference SAI and the neighboring SAIs. Thanks to the collaborative nature of the proposed filter, the reference and neighboring SAIs are thus denoised at once. A second iteration is usually necessary to process the few remaining noisy pixels in the neighboring SAIs.

Once the current reference and neighboring SAIs are fully processed, a new reference SAI is selected among the SAIs remaining to denoise. The new reference SAI is chosen such that it has the most neighboring SAIs remaining to denoise. The process iterates until all SAIs in L^{noisy} are denoised. This main loop is used for both hard-thresholding and Wiener filtering steps, and produces the basic estimate L^{basic} and final estimate L^{final} respectively.

In the next sections, we explain the different steps of the filter in details.

B. Construction of 5D patches

For each 2D patch P of size $k \times k$ in the reference SAI I_{ref} , a 5D patch is created by exploiting redundancies within I_{ref} and in-between I_{ref} and its neighboring SAIs $\{I_{s,t}\}$, $(s, t) \in [1, n_a] \times [1, n_a]$. Note that to reduce the processing time, the loop over the overlapping 2D patches in a SAI is done with a step of p pixels in row and column.

1) *4D patch*: First, a 4D patch is created by finding in each neighboring SAI the 2D patch closest to P (see Fig. 2), which can be assimilated to a disparity compensation step using a block matching algorithm. Formally, the 4D patch of size $n_a \times n_a \times k \times k$ is defined as:

$$\mathcal{P}_{disp}(P) = \{Q_{disp}^{s,t} : Q_{disp}^{s,t} = \arg \min_{Q_{s,t}} d(P, Q_{s,t}), \\ Q_{s,t} \in SW_{disp}^{s,t}, \\ d(P, Q_{disp}^{s,t}) \leq \tau_{disp}\} \quad (2)$$

where $d(P, Q)$ is the normalized quadratic distance between patches, $SW_{disp}^{s,t}$ is a search window in $I_{s,t}$ of size $n_{disp} \times n_{disp}$ centered on the position of P (see Fig. 2), and τ_{disp} is the distance threshold for d under which patches are assumed similar.

The goal of the threshold τ_{disp} is to discard patches which are not similar enough to the reference patch, and thus to be robust to occlusions. Together with the disparity compensation, it ensures the smoothness and homogeneity of 2D patches taken along the angular dimensions (also called 2D angular patches, see Fig. 2).

2) *5th dimension - self-similarities*: We then search for the set of patches similar to P in the reference SAI, defined by:

$$\mathcal{P}_{sim}(P) = \{Q_{sim} : d(P, Q_{sim}) \leq \tau_{sim}, Q_{sim} \in SW_{sim}\} \quad (3)$$

where τ_{sim} is the distance threshold for d under which patches are assumed similar, and SW_{sim} is a search window in I_{ref} of size $n_{sim} \times n_{sim}$ centered on the position of P (see Fig. 2). A maximum of N patches are retained in $\mathcal{P}_{sim}(P)$ in order to limit the complexity.

Finally, the 5D patch of size $n_a \times n_a \times k \times k \times N$, denoted $\mathbb{P}(P)$, is then built by stacking along a 5th dimension the 4D patches built from all patches similar to the reference patch (see Fig. 2):

$$\mathbb{P}(P) = \{\mathcal{P}_{disp}(Q) : Q \in \mathcal{P}_{sim}(P)\} \quad (4)$$

Note that for the Wiener filtering step, a first 5D patch $\mathbb{P}^{basic}(P)$ is obtained from the basic estimate L^{basic} , and a second 5D patch $\mathbb{P}(P)$ is then created by taking the collocated patches in L^{noisy} .

C. Collaborative filtering in the 5D transform domain

The collaborative property of the filter means that when a 2D patch P is processed, all the 2D patches constituting the 5D patch $\mathbb{P}(P)$ are filtered together. The filter acts in the 5D transform domain, and the filtered patches are obtained after applying the inverse transform.

In the first step, hard-thresholding is applied on the 5D transform coefficients of the patch:

$$\mathbb{P}(P)^{hard} = \tau_{5D}^{hard^{-1}}(\gamma(\tau_{5D}^{hard}(\mathbb{P}(P)))) \quad (5)$$

where τ_{5D}^{hard} is a 5D transform, γ is a hard-thresholding operator with threshold $\lambda_{5D}^{hard}\sigma$ and σ^2 is the variance of the AWGN:

$$\gamma(x) = \begin{cases} 0 & \text{if } |x| \leq \lambda_{5D}^{hard}\sigma \\ x & \text{otherwise} \end{cases} \quad (6)$$

For the second step, the hard-thresholding is replaced by a Wiener filter, using the 5D patch $\mathbb{P}^{basic}(P)$ as a guide. The empirical Wiener coefficients are defined by:

$$\omega_P(\xi) = \frac{|\tau_{5D}^{wien}(\mathbb{P}^{basic}(P))(\xi)|^2}{|\tau_{5D}^{wien}(\mathbb{P}^{basic}(P))(\xi)|^2 + \sigma^2} \quad (7)$$

where τ_{5D}^{wien} is a 5D transform, and σ^2 is the variance of the AWGN.

The Wiener collaborative filtering of $\mathbb{P}(P)$ is then performed as the element-wise multiplication of the 5D transform of the noisy light field $\tau_{5D}^{wien}(\mathbb{P}(P))$ with the Wiener coefficients ω_P . An estimate of the 5D patch is obtained as:

$$\mathbb{P}^{wien}(P) = \tau_{5D}^{wien^{-1}}(\omega_P \cdot \tau_{5D}^{wien}(\mathbb{P}(P))) \quad (8)$$

The 5D transform consists in practice of three cascaded transforms: a 2D transform denoted τ_{2D}^s is applied on the spatial dimensions of the 5D patch, followed by a 2D transform denoted τ_{2D}^a applied on the angular dimensions (see Fig. 3). Finally, a 1D transform denoted τ_{1D} is applied on the 5th dimension of the 5D patch. The choice of these transforms will be discussed in section III-E.

D. Aggregation

Once the collaborative filtering is done, we get an estimate for each 2D patch belonging to a 5D patch, and thus multiple estimates are obtained for each pixel. The final pixel estimate is obtained as a weighted average of all estimates:

$$L^{step}(x) = \frac{\sum_P w_P^{step} \sum_{Q \in \mathbb{P}(P)} \chi_Q(x) L_{Q,P}^{step}(x)}{\sum_P w_P^{step} \sum_{Q \in \mathbb{P}(P)} \chi_Q(x)} \quad (9)$$

where *step* can be either of the two steps, denoted *hard* or *wiener*, $L_{Q,P}^{step}(x)$ is the estimate of the value of the pixel x belonging to the 2D patch Q obtained during collaborative filtering of the reference patch P , and $\chi_Q(x) = 1$ if and only if $x \in Q$, 0 otherwise.

For the hard-thresholding step, the aggregation weights are obtained as:

$$w_P^{hard} = \begin{cases} (N_P^{hard})^{-1} & \text{if } N_P^{hard} \geq 1 \\ 1 & \text{otherwise} \end{cases} \quad (10)$$

where N_P^{hard} is the number of non-zero coefficient in the 5D patch after hard-thresholding: $\gamma(\tau_{5D}^{hard}(\mathbb{P}(P)))$.

For the Wiener filtering step, the aggregation weights are computed as $w_P^{wien} = \|\omega_P\|_2^{-2}$.

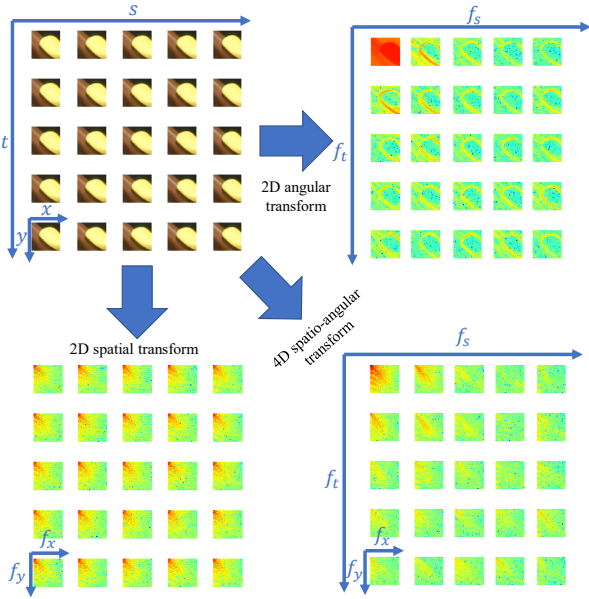


Fig. 3. Example of a 4D transform spectrum (bottom right) obtained for a 4D patch (top left) taken from the LegoKnights light field. The 4D transform is obtained by cascading a 2D spatial transform (bottom left) and a 2D angular transform (top right). In this example the 2D DCT was used for both transforms. Notice that the 4D transform spectrum is sparser than the individual 2D transform spectra.

E. Choice of the 5D transform

As indicated in section III-C, the 5D transform is a combination of 2D transforms τ_{2D}^s and τ_{2D}^a applied on the spatial and the angular dimensions respectively, and a 1D transform τ_{1D} applied on the 5th dimension corresponding to similarities. The choice of these transforms has a strong impact on the denoising performances and should be considered carefully.

For the choice of τ_{2D}^s and τ_{1D} , we rely on the analysis carried out in [17]. The 2D spatial transform can either be a normalized 2D DCT, or a bi-orthogonal spline wavelet denoted Bior1.5, where the vanishing moments of the decomposing and reconstructing wavelet functions are 1 and 5 respectively. Better performance is achieved when the transform is different for the hard-thresholding step and for the Wiener filtering step. For the 1D transform, we choose the Haar wavelet.

One of the main advantages of the proposed filter compared to existing work is that we can fully exploit the 2D angular redundancies by applying τ_{2D}^a over the angular dimensions of the 5D patches. If we look at a 5D patch from the perspective of its angular dimensions, we can consider that it consists of $k \times k \times N$ 2D angular patches of size $n_a \times n_a$. Thanks to the disparity compensated creation of the 5D patch, the 2D angular patches are smooth and we thus obtain a sparse representation in the 2D angular transform domain. By cascading the 2D angular and spatial transforms, we finally obtain very sparse representations of 4D patches, which are further compacted when the 1D transform is applied over the 5th dimension. We show in Fig. 3 an example of 2D angular and spatial transform spectra obtained from a 4D patch. The 4D transform spectrum obtained from the combination of the 2D transforms is clearly

sparser than the 2D transform spectra.

Note that the 2D angular patch smoothness is further enforced by the threshold τ_{disp} introduced in section III-B to be robust to occlusions. However, this implies that some 2D angular patches may have a non-square shape. Thus, we propose to use for τ_{2D}^a the 2D shape adaptive DCT (SADCT). This transform was first introduced for coding arbitrarily shaped segment of images [21], and was later used for image denoising [22]. The 2D SADCT can thus handle non-square patches by applying 1D DCT transform of varying length on the columns and then the rows of the patch.

F. Color light fields

As for the BM3D algorithm, color light fields are processed by first transforming the noisy RGB SAIs to a luminance-chrominance space. The 5D patches are then created using the luminance channel only. The collaborative filtering and aggregation steps are then applied on each channel separately. Finally, the denoised SAIs are returned to the RGB space. We use in this paper the opponent color transformation (OPP) proposed in [16].

IV. RESULTS

We discuss in this section the performance of the proposed approach against relevant state-of-the-art methods. We implemented our algorithm in C/C++ based on the BM3D implementation of [17]¹. Code and detailed results will be available online². We choose to compare to the hyperfan filter applied on the global 4D transform domain [4], denoted HF4D, independent denoising of the SAIs using the BM3D filter [16], the two-step EPI denoising of [9] using the BM3D filter, and the VBM4D filter applied on sequence of SAIs or EPIs [12].

A. Synthetic noise

We first study the denoising performances in the case of synthetic AWGN. The experiments are conducted on a dataset of twelve light fields from EPFL captured with a Lytro Illum camera [23] and a second dataset of twelve light fields from Stanford captured with a gantry [5]. For the EPFL dataset, the 15 \times 15 SAIs of resolution 434 \times 625 were extracted using Dansereau's Matlab light field toolbox [24]. Note that for the Stanford dataset, the spatial resolution of the SAIs was cropped to 512 \times 512 pixels to speed up the experiments, but the angular resolution of 17 \times 17 was maintained.

We set five noise levels, $\sigma = 10, 20, 30, 40, 50$, and the parameters of our method are shown in Table I. For the BM3D and VBM4D filters we use standard parameter values [17] [20]. Additional parameters related to our method are tuned empirically.

We evaluate the denoising performances using the PSNR. For each light field the average PSNR is computed over all SAIs. We give in Table II the average results for both the EPFL dataset and the Stanford dataset. The line Δ PSNR corresponds

¹<http://www.ipol.im/pub/art/2012/l-bm3d/>

²<https://v-sense.scss.tcd.ie/?p=893>

TABLE I
LFBM5D MAIN PARAMETERS

	EPFL (Lytro Illum)		Stanford (Gantry)	
	<i>hard</i>	<i>wiener</i>	<i>hard</i>	<i>wiener</i>
n_a	3			
k	16	8	16	8
N	1	8	8	16
p	3	3	4	4
n_{disp}	7	7	17	17
n_{sim}	37			
τ_{disp} if $\sigma \leq 35$	3000	2000	3000	2000
τ_{disp} otherwise	5000			
τ_{sim} if $\sigma \leq 35$	3000	2000	3000	2000
τ_{sim} otherwise	5000			
χ_{5D}^{hard}	2.7		2.7	
τ_{2D}^s	2D Bior1.5	2D DCT	None	2D DCT
τ_{2D}^a		2D SADCT		
τ_{1D}		Haar		

TABLE II
AVERAGE DENOISING PERFORMANCES IN PSNR

Method	$\sigma=10$	$\sigma=20$	$\sigma=30$	$\sigma=40$	$\sigma=50$
	EPFL dataset (Lytro Illum)				
HF4D [11]	31.070	25.798	22.607	20.338	18.586
BM3D [16]	35.421	32.852	31.357	30.247	29.321
BM3D EPI [9]	36.088	33.476	31.905	30.712	29.671
VBM4D [20]	36.075	33.522	31.923	30.674	29.630
VBM4D EPI [12]	36.129	33.510	31.925	30.719	29.721
LFBM5D 1st step	34.388	32.810	31.684	30.743	29.911
LFBM5D 2nd step	36.503	34.214	32.868	31.843	30.987
Δ PSNR	0.374	0.692	0.943	1.124	1.266
	Stanford dataset (Gantry)				
HF4D [11]	30.577	25.432	22.213	19.921	18.164
BM3D [16]	38.805	35.268	33.126	31.560	30.267
BM3D EPI [9]	38.895	35.645	33.489	31.896	30.594
VBM4D [20]	39.269	35.588	33.212	31.436	30.019
VBM4D EPI [12]	38.697	35.604	33.558	32.026	30.809
LFBM5D 1st step	39.340	35.817	33.377	31.496	29.971
LFBM5D 2nd step	40.389	37.772	36.031	34.671	33.511
Δ PSNR	1.120	2.127	2.473	2.645	2.701

to the PSNR gap between the proposed approach and the best state-of-the-art method.

These results show a clear progression of the denoising performance which increases with the extent to which the light field structure is taken into account. Our approach clearly outperforms the state-of-the-art methods for all noise levels, with a notably significant improvement when the noise level increases.

B. Lenslet camera noise removal

In the previous section, we added synthetic noise to light fields in order to compute objective measures where the input light field is considered as ground truth. However, light fields captured with a lenslet camera such as the Lytro Illum already exhibit camera noise. In this section, we apply our approach to the light fields from the EPFL dataset in order to remove such noise.

Visual results are displayed in Fig. 5 for three light fields from the EPFL dataset. In the top row, we show one SAI from each input light field, where the camera noise patterns and dead pixels are clearly visible. In the bottom row, we present the results after applying our denoising method, which clearly reduces the camera noise and removes the dead pixels while preserving fine details and edges.




Fig. 4. Visual results of denoising with AWGN on the center SAI of a light field. (Best viewed in color and zoomed)

TABLE III
AVERAGED ESTIMATED NOISE LEVEL [25] FOR THE EPFL DATASET

Method	Input	VBM4D EPI	VBM4D	LFBM5D
$\sigma_{est} * 100$	84.1	28.3	6.9	6.0

In order to quantify the noise reduction, we perform blind noise level estimation [25] before and after denoising. We compare our own approach against the two best state-of-the-art methods, the VBM4D filter and the VBM4D applied on EPIs [12]. We give in Table III the estimated values of the noise variance σ_{est} averaged over the dataset, which demonstrate that our method performs a better noise reduction than state-of-the-art methods.

V. CONCLUSION

In this paper, we introduced the LFBM5D filter, which extends the principles of the BM3D filter to light field denoising. By creating disparity compensated 4D patches which are stacked together with similar 4D patches along a 5th dimension, we obtain in the 5D transform a very sparse representation where the noise is well separated from the underlying signal and can be efficiently filtered. The 5D transform is obtained by cascading a 2D spatial transform, a 2D angular transform, and a 1D transform over the 5th dimension. We propose to use the SA-DCT for the 2D angular transform in order to be robust to occlusion and enforce the sparsity in the transform domain. As our method fully exploits the 4D structure of the light field and the self-similarities existing in natural images, we are able to outperform the states-of-the-art methods for light fields captured with different camera configurations.

In future work we intend to exploit the disparity values estimated during the creation of the disparity compensated 4D

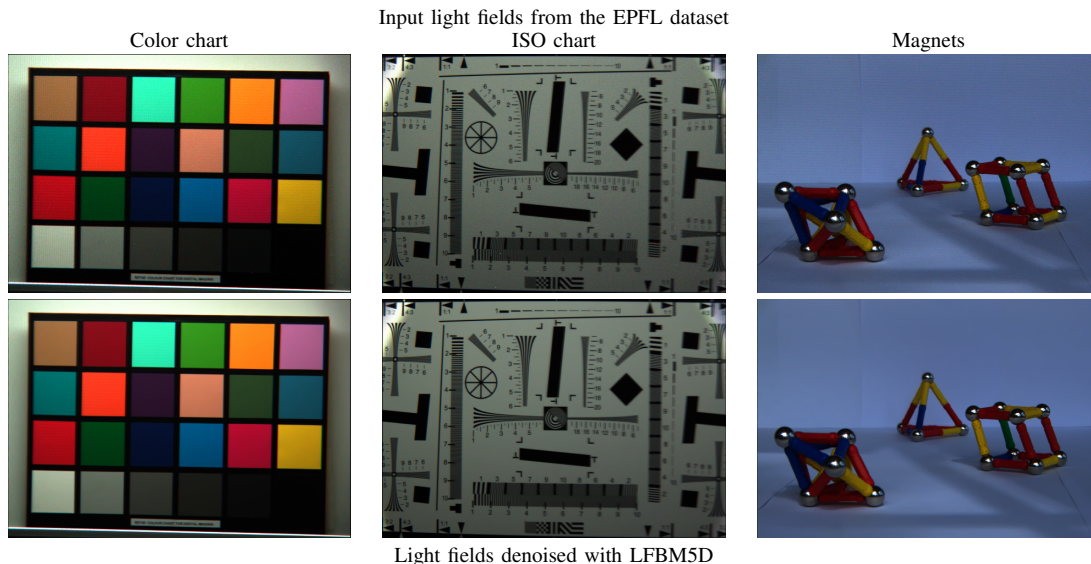


Fig. 5. Visual results of LFBM5D denoising (bottom row) applied on light fields captured with the Lytro Illum camera [23] (top row). Notice that the overall camera noise pattern, especially visible on light fields captured in low light conditions such as Magnet, is considerably reduced, while preserving high frequency and fine details, e.g. on ISO chart. Dead pixels, especially visible in Color and ISO charts, are also removed. (**Best viewed in color and zoomed**)

patches, in order to create depth maps from noisy light fields, which is a challenging task. The choice of the transforms used to create the 5D transform is crucial, and we plan to explore other 2D angular transforms, such as edge-avoiding wavelets or graph Fourier transforms [26]. Finally, the improved denoising performance of our method comes at the price of an increased complexity. We also aim to study in details the complexity in order to reduce it.

ACKNOWLEDGMENT

This work was supported by TCHPC (Research IT, Trinity College Dublin). All calculations were performed on the Lonsdale cluster maintained by the Trinity Centre for High Performance Computing. This cluster was funded through grants from Science Foundation Ireland.

REFERENCES

- [1] M. Levoy and P. Hanrahan, "Light field rendering," in *Proc. SIGGRAPH*, 1996, pp. 31–42.
- [2] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan, "Light Field Photography with a Hand-Held Plenoptic Camera," CSTR 2005-02, Tech. Rep., Apr. 2005.
- [3] A. Isaksen, L. McMillan, and S. J. Gortler, "Dynamically reparameterized light fields," in *Proc. SIGGRAPH*, 2000, pp. 297–306.
- [4] B. Goldluecke, "Globally consistent depth labeling of 4d light fields," in *Proc. CVPR*, 2012, pp. 41–48.
- [5] "The stanford light field archive," <http://lightfield.stanford.edu/lfs.html>, accessed: 29-05-2017.
- [6] B. Wilburn, N. Joshi, V. Vaish, E.-V. Talvala, E. Antunez, A. Barth, A. Adams, M. Horowitz, and M. Levoy, "High performance imaging using large camera arrays," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 765–776, Jul. 2005.
- [7] "The lytro illum camera," <https://www.lytro.com/imaging>, accessed: 23-05-2017.
- [8] K. Mitra and A. Veeraraghavan, "Light field denoising, light field superresolution and stereo camera based refocussing using a GMM light field patch prior," in *Proc. CVPR Workshops*, jun 2012, pp. 22–28.
- [9] Z. Li, H. Baker, and R. Bajcsy, "Joint image denoising using light-field data," *Proc. ICMEW*, 2013.
- [10] B. Goldluecke and S. Wanner, "The variational structure of disparity and regularization of 4D light fields," *Proc. CVPR*, pp. 1003–1010, 2013.
- [11] D. G. Dansereau, D. L. Bongiorno, O. Pizarro, and S. B. Williams, "Light field image denoising using a linear 4D frequency-hyperfan all-in-focus filter," in *Proc. SPIE*, vol. 8657, Feb 2013, p. 86570P.
- [12] A. Sepas-Moghaddam, P. L. Correia, and F. Pereira, "Light field denoising: exploiting the redundancy of an epipolar sequence representation," in *Proc. 3DTV-Con*, Jul 2016, pp. 1–4.
- [13] D. L. Donoho, "De-noising by soft thresholding," *Symp. Wavelet Theory*, vol. 76, p. 1, 1992.
- [14] S. P. Ghael, A. M. Sayeed, and R. G. Baraniuk, "Improved wavelet denoising via empirical Wiener filtering," in *Proc. SPIE*, Oct 1997, pp. 389–399.
- [15] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Proc. CVPR*, 2005, pp. 60–65.
- [16] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering," *IEEE Trans. on Image Processing*, vol. 16, no. 8, pp. 2080–2095, 2007.
- [17] M. Lebrun, "An Analysis and Implementation of the BM3D Image Denoising Method," *Ipol.Im*, 2012.
- [18] L. Shao, R. Yan, X. Li, and Y. Liu, "From heuristic optimization to dictionary learning: A review and comprehensive comparison of image denoising algorithms," *IEEE Trans. on Cybernetics*, vol. 44, no. 7, pp. 1001–1013, 2014.
- [19] P. Jain and V. Tyagi, "A survey of edge-preserving image denoising methods," *Information Systems Frontiers*, vol. 18, no. 1, pp. 159–170, 2016.
- [20] M. Maggioni, G. Boracchi, A. Foi, and K. Egiazarian, "Video denoising, deblocking, and enhancement through separable 4-D nonlocal spatiotemporal transforms," *IEEE Trans. on Image Processing*, vol. 21, no. 9, pp. 3952–3966, 2012.
- [21] T. Sikora, "Low complexity shape-adaptive DCT for coding of arbitrarily shaped image segments," *Signal Processing: Image Communication*, vol. 7, no. 4-6, pp. 381–395, 1995.
- [22] A. Foi, V. Katkovnik, and K. Egiazarian, "Pointwise shape-adaptive DCT for high-quality denoising and deblocking of grayscale and color images," *IEEE Trans. on Image Processing*, vol. 16, no. 5, pp. 1395–1411, May 2007.
- [23] M. Rerabek and T. Ebrahimi, "New Light Field Image Dataset," in *Proc. QoMEX*, 2016.
- [24] D. G. Dansereau, O. Pizarro, and S. B. Williams, "Decoding, calibration and rectification for lenselet-based plenoptic cameras," in *Proc. CVPR*, 2013, pp. 1027–1034.
- [25] X. Liu, M. Tanaka, and M. Okutomi, "Single-image noise level estimation for blind denoising," *IEEE Trans. on Image Processing*, vol. 22, no. 12, pp. 5226–5237, Dec. 2013.
- [26] W. Hu, G. Cheung, X. Li, and O. C. Au, "Graph-based joint denoising and super-resolution of generalized piecewise smooth images," in *Proc. ICIP*, Oct 2014, pp. 2056–2060.