Paper 22-3 has been designated as a Distinguished Student Paper at Display Week 2025. The full-length version of this paper appears in a Special Section of the *Journal of the Society for Information Display (JSID)* devoted to Display Week 2025 Distinguished Papers. This Special Section will be freely accessible until December 31, 2025 via:

https://sid.onlinelibrary.wiley.com/doi/full/10.1002/jsid.2075

Authors that wish to refer to this work are advised to cite the full-length version by referring to its DOI:

https://doi.org/10.1002/jsid.2075

BDLUT: Blind Image Denoising with Hardware-Optimized Look-Up Tables

Boyu Li[1,*], Zhilin Ai[1,*], Baizhou Jiang[1], Binxiao Huang[1], Jason Chun Lok Li[1], Jie Liu[2], Zhengyuan Tu[2], Guoyu Wang[2], Daihai Yu[2], Ngai Wong[1]

[1]Department of Electrical and Electronic Engineering, The University of Hong Kong, Hong Kong

[2]TCL Corporate Research (HK) Co., Ltd. Hong Kong * *Equal contributors*

Abstract

We propose BDLUT(-D) that novelly combines blind denoising with hardware-optimized lookup tables (LUTs) for resource-efficient edge computing. While BDLUT describes the LUT-based network architecture, BDLUT-D represents BDLUT trained with a specialized noise degradation model. BDLUT-D achieves superior denoising, outperforming state-of-the-art (SOTA) LUT methods by up to 2.42 dB on mixed-noise-intensity benchmark datasets using only 66KB storage. Our FPGA implementation demonstrates significant hardware efficiency, reducing logic resource consumption by over 10x and storage requirements by 75% compared to DNN accelerators, while achieving 57% higher speed than bilateral filtering methods.

Author Keywords

Blind Denoising, Lookup Table, Algorithm-Hardware Co-Design

1. Introduction

Image denoising remains a fundamental challenge in computer vision, particularly for edge devices with limited computational resources. Traditional methods like bilateral filtering (BF) [1–3] and block-matching (CBM3D) [4] often struggle with complex noise patterns, while deep neural networks (DNNs) [5, 6] require substantial computational resources despite their superior performance. This creates a critical gap between denoising quality and hardware efficiency.

Real-world image noise often deviates significantly from idealized models, exhibiting complex characteristics that are spatially dependent and non-Gaussian. While Blind super-resolution (SR) techniques[7], and by extension blind denoising can handle unknown noise distributions, existing solutions typically rely on computationally intensive DNNs. Look-up tables (LUTs) [8–11] especially Hundred-Kilobyte Lookup Tables (HKLUT) [12] have emerged as a promising approach for efficient image processing, but their application to blind denoising remains unexplored.

This work introduces BDLUT, integrating blind denoising with hardware-optimized LUTs to address both noise uncertainty and computational efficiency. As shown in Fig. 1, BDLUT achieves exceptional denoising performance while maintaining minimal computational overhead. Our key contributions include:

- First integration of blind denoising with LUT-based processing, validated through extensive experiments on adequate benchmarks.

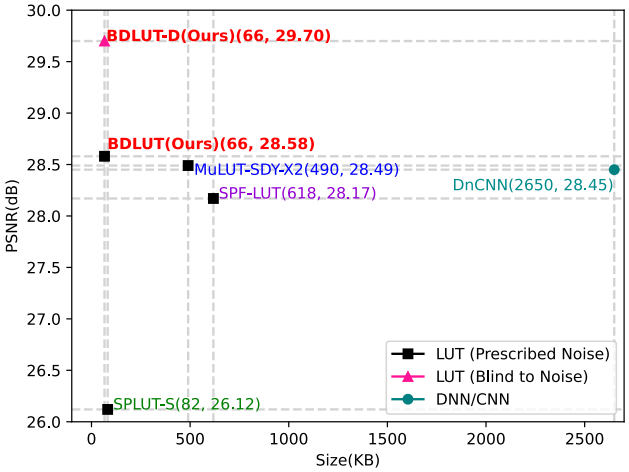

Figure 1: Comparison of PSNR (Test set: Composed of one-third of the images from Kodak24, each corrupted by AWGN with noise levels of $\sigma = 15$, $\sigma = 25$, and $\sigma = 50$. Train set: DIV2K with $\sigma = 15$ AWGN for non-blind LUT) and model size between the proposed method and other methods.

- Enhanced kernel design based on HKLUT for improved denoising performance and hardware efficiency, achieving state-of-the-art (SOTA) results with only 66KB storage
- Efficient FPGA implementation demonstrating significant resource savings compared to conventional DNN accelerators

2. Development of BDLUT

2.1 Optimized Kernel Design

Building upon HKLUT's success in super-resolution [12], we develop an optimized kernel design for blind denoising. The final prediction $\hat{y}_i$ is computed as:

$$\hat{y}_i = \frac{1}{N}\frac{1}{M_k}\sum_{k=0}^{N}\sum_{j=0}^{M_k} R_j^{-1}(LUT_k(R_j(x_i))) \tag{1}$$

where x_i represents the noisy input, LUT_k is the k^{th} LUT, R_j denotes the j^{th} rotation operation, N is the kernel count, and M_k is the rotation count per kernel.

Our design employs four types of 3-pixel kernels with rotation ensemble to effectively cover a 5×5 receptive field, maintaining 8 overlapping pixels around the pivot. This configuration balances

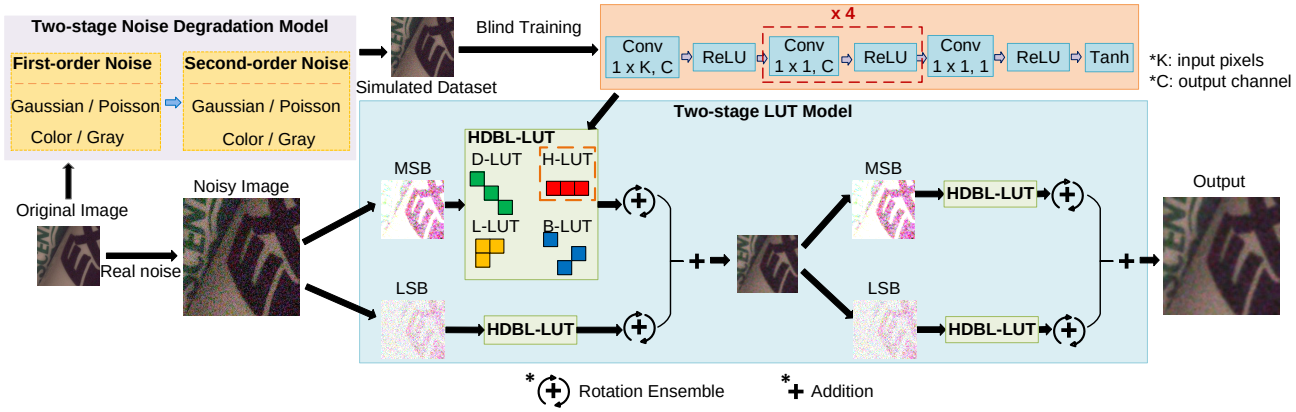

Figure 2: Architecture of BDLUT, showing the two-stage noise degradation model for blind training and the dual-branch LUT structure with rotation ensemble.

denoising performance with hardware efficiency. The kernel structure and CNN training architecture are illustrated in Fig. 2. Notably, we utilize parallel MSB and LSB branches with identical RF size and kernel numbers, maximizing denoising effectiveness while maintaining reasonable model size.

2.2 Real-World Blind Noise Degradation Model

To handle real-world noise effectively, we develop a two-stage stochastic degradation model extending [7]. The model incorporates:

- Gaussian additive noise: Applied independently to RGB channels (color noise) or uniformly (gray noise)
- Poisson noise: Intensity-dependent noise following Poisson distribution is applied as color/gray noise

As shown in Fig. 2, noise is applied in two stages with randomly determined parameters. During each stage, noise type and application method are probabilistically selected, with intensity parameters randomly sampled within predefined ranges. This approach effectively simulates the diverse characteristics of real-world image noise.

3. Hardware Implementation

Traditional hardware denoising implementations often rely on bilateral filtering [1–3] due to its simplicity, despite limitations in quality and resource efficiency. While DNN-based approaches like Light-DnCNN [13] offer improved quality, they require significant computational resources.

Our work presents the first FPGA implementation of a LUT-based denoising module. The architecture, illustrated in Fig. 3, processes 5×5 RGB pixel blocks through:

- Channel separation for parallel processing of R, G, B components
- Four kernel types (H, D, B, L) with 0°, 90°, 180°, and 270° rotations
- Efficient memory organization using Block RAM (BRAM)
- Hardware-optimized address generation and result averaging

This design eliminates multiplication operations while maintaining high throughput, enabling efficient real-time processing of video streams.

4. Experiments

4.1 Experiment Setup

Datasets and Training: We train on DIV2K[16] using two protocols: (1) AWGN with $\sigma = 15, 25, 50$ for standard denoising, and (2) our blind noise model for real-world denoising. The blind noise model combines Gaussian and Poisson noise (probability 0.5 each), with noise parameters randomly sampled: $\sigma \in [0, 50]$ for Gaussian noise and scale $\in [0.05, 3]$ for Poisson noise (reduced to $[0.05, 0.25]$ for second-stage noise). Gray noise probability is set to 0.4 for both stages.

Implementation Details: Training runs for 200k iterations with batch size 16 on NVIDIA RTX 3090 GPUs using Adam optimizer ($\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 1e-8$) and MSE loss. Learning rate starts at 5×10^{-4}, decaying by 0.1 at 100k and 150k iterations. The network is then converted to LUT format for deployment. The LUT PE implementation targets ZCU102 FPGA, with COE files preloaded into BRAM.

Evaluation: We evaluate on standard benchmarks including Set12[5], BSD68[17], CBSD68[17], Kodak24[18], McMaster[19], Urban100[20] for synthetic noise, and SIDD[21] for real-world noise. The mixed-noise-intensity dataset is constructed by adding AWGN at σ =15, 25, and 50 to three equal dataset partitions of aforementioned datasets (excluding SIDD). We compare against SOTA methods including classical (BM3D[4], CBM3D[4], WNNM[15], MC-WNNM[14]), DNN-based (DnCNN[5], SwinIR[6]), and LUT-based approaches (SR-LUT[8], MuLUT[9], SPF-LUT[11]).

4.2 Performance Analysis

Denoising Quality: In Tables 1, 3, BDLUT achieves competitive performance and requires the least storage. In Table 5, BDLUT-D shows practical advantages by adapting to diverse unknown noise conditions. However, its performance is limited by the discrepancy between its noise-agnostic training and the noise characteristics in non-blind evaluation settings.

General AWGN denoising performance: In Table 4 BDLUT-D achieves the best results, demonstrating its superior generalization

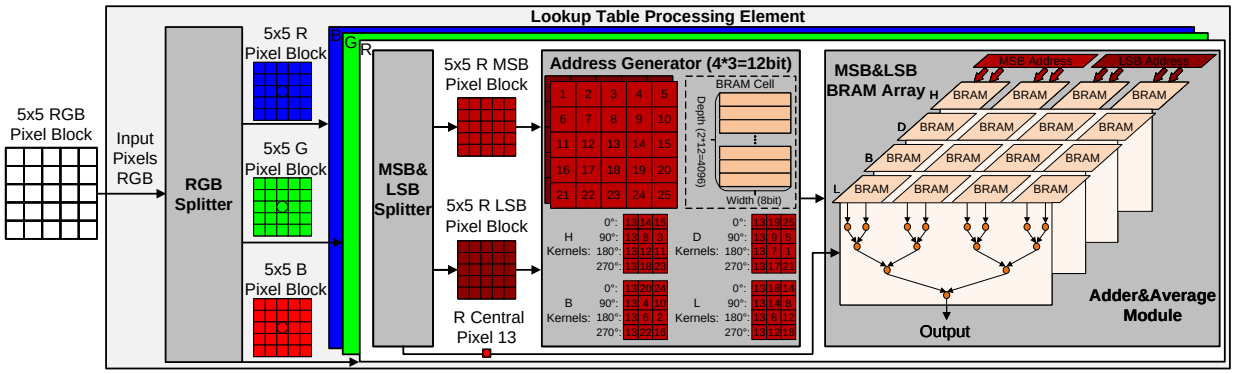

Figure 3: Architecture of the LUT Processing Element (PE), showing the RGB channel separation, parallel processing paths, and efficient memory organization.

Table 1: Size and Quantitative comparison (PSNR/dB) for color image denoising on 4 fixed-noise-intensity benchmark datasets. Smallest size of each model is shown in bold.

Category	Method	Size	CBSD68			Kodak24			Urban100			McMaster		
			$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$
LUT-based	SR-LUT[8]	82KB	32.72	29.60	25.41	33.44	30.31	25.98	32.42	28.69	24.05	34.69	31.03	26.11
	MuLUT-SDY-X2[9]	490KB	33.20	30.48	27.12	34.01	31.32	28.04	33.23	30.00	25.84	35.56	32.77	29.18
	SPF-LUT-Net[11]	618KB	34.02	31.30	28.03	34.85	32.23	29.02	34.49	31.30	27.29	36.52	33.78	30.25
	BDLUT-D(Ours)	**66KB**	31.46	30.52	25.02	32.05	31.19	25.88	30.62	29.71	25.30	34.08	32.80	26.60
	BDLUT(Ours)	**66KB**	33.41	30.56	27.25	34.18	31.42	27.93	33.1	29.92	26.2	35.67	32.84	29.12
Classical	CBM3D[4]	—	33.52	30.71	27.38	34.28	31.68	28.46	33.92	31.35	27.94	34.06	31.66	28.51
	MC-WNNM[14]	—	29.61	26.73	23.18	30.04	26.34	21.18	30.06	27.01	23.01	30.65	27.74	23.96
DNN	DnCNN[5]	2.65MB	33.89	31.23	28.01	34.48	32.03	28.85	32.98	30.81	27.59	33.45	31.52	28.62
	SwinIR[6]	117MB	34.42	31.78	28.56	35.34	32.89	29.79	35.13	32.90	29.82	35.61	33.20	30.22

Table 2: Ablation study results (PSNR) showing impact of different kernel configurations and receptive field (RF) sizes for MSB and LSB branches. Best performance shown in bold.

Method	RF	Size	CBSD68			Kodak24			Urban100			McMaster		
MSB / LSB	MSB / LSB	(KB)	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$
HD / HD	3×3 / 3×3	3	32.76	29.57	26.42	33.42	30.14	26.96	31.84	28.46	24.75	34.19	31.63	28.16
HD / HDB	3×3 / 5×5	26.2	32.94	29.96	26.47	33.60	30.62	27.02	32.46	28.97	24.69	35.00	32.03	28.26
HDB / HD	5×5 / 3×3	26.2	33.09	30.01	26.52	33.82	30.79	27.17	32.64	28.97	24.97	35.28	31.49	27.77
HDB / HDB	5×5 / 5×5	49.5	33.35	30.50	26.96	34.11	31.32	27.59	33.04	29.86	25.43	35.66	32.74	28.71
HDBL / HDBL	5×5 / 5×5	66	**33.41**	**30.56**	**27.25**	**34.18**	**31.42**	**27.93**	**33.17**	**29.92**	**25.94**	**35.67**	**32.84**	**29.12**

Table 3: Quantitative comparison(PSNR) for grayscale image denoising on benchmark datasets.

Category	Method	BSD68			Set12		
		$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$
LUT-based	SR-LUT[8]	29.78	26.85	22.39	30.42	27.19	22.62
	MuLUT-SDY-X2[9]	30.63	28.18	24.97	31.50	28.94	25.46
	SPF-LUT[11]	31.17	28.50	25.59	32.11	29.47	26.21
	BDLUT-D(Ours)	29.61	27.62	23.39	30.51	28.37	23.69
	BDLUT(Ours)	30.41	27.81	24.43	31.20	28.55	24.89
Classical	BM3D[4]	31.07	28.57	25.62	32.37	29.97	26.72
	WNNM[15]	31.37	28.83	25.87	32.70	30.26	27.05
DNN	DnCNN[5]	31.73	29.23	26.23	32.86	30.44	27.18
	SwinIR[6]	31.97	29.50	26.58	33.36	31.01	27.91

and robustness. BDLUT-D's adaptive degradation model effectively handles diverse and dynamic real-world noises through its flexible parameter settings.

Ablation Study: Table 2 validates our kernel design choices. The HDBL/HDBL configuration provides an optimal trade-off between performance and efficiency. The results show balanced importance between MSB and LSB processing, with PSNR differences within 0.2 dB between different configurations.

Hardware Efficiency: Our FPGA implementation achieves 375 MHz operation frequency with minimal resource utilization: 7137 LUTs, 162 FFs, and 3026 KB BRAM, requiring no DSP units. As shown in Table 6, BDLUT significantly outperforms existing hardware implementations, achieving 57% higher processing speed than

Table 4: Quantitative comparison (PSNR/dB) for color and grayscale image denoising on mixed-noise-intensity benchmark datasets which are sample one-third of images from each AWGN-corrupted test set (σ = 15, 25, 50). Best performance is shown in bold.

Method	Train set (σ)	CBSD68	Kodak24	Urban100	McMaster	BSD68	Set12
SR-LUT[8]	15	28.26	28.12	27.78	29.15	23.85	24.17
	25	28.13	28.70	27.44	29.71	24.94	25.54
	50	27.22	28.22	26.12	29.20	24.88	25.64
MuLUT-SDY-X2[9]	15	28.23	28.49	28.48	29.76	24.24	24.67
	25	28.11	28.72	27.95	30.15	25.11	25.98
	50	27.21	28.04	26.40	29.23	25.21	25.97
SPF-LUT-Net[11]	15	27.95	28.17	28.24	29.54	23.85	24.25
	25	27.94	28.41	28.47	30.27	25.13	26.08
	50	27.57	28.69	27.14	29.90	25.74	26.94
BDLUT (Ours)	15	28.26	28.58	28.12	29.46	24.21	24.56
	25	28.46	29.19	28.17	30.49	25.45	26.31
	50	26.82	27.93	26.42	29.75	25.48	26.41
BDLUT-D (Ours)	-	**29.00**	**29.70**	**28.54**	**31.16**	**26.34**	**27.18**

Table 5: Quantitative comparison (PSNR/dB) for real-world image denoising on validation datasets.

Dataset	Method				
	SRLUT[8]	MuLUT-SDY-X2[9]	SPF-LUT-Net[11]	BDLUT-D	BDLUT
SIDD	34.30	40.09	40.20	36.44	39.37

Table 6: Hardware implementation comparison across FPGA platforms, showing resource utilization and performance metrics for 1024×1024 image processing.

Design	FPGA Platform	Fmax (MHz)	Frame Rate	Resource Utilization			
				LUT	FF	BRAM (Kb)	DSP
[1]	XC7Z020	238	226.7	**1357**	2118	72	32
[2]	XC7Z020	250	52.45	1594	2399	**0**	6
[3]	Virtex-5	242	230	2529	1917	144	27
[13]	ZCU106	200	1.25	159967	29851	11232	736
Ours	ZCU102	**375**	**357.7**	7137	**162**	3026	**0**

bilateral filtering methods while consuming only 4.5% of the logic resources required by DNN-based accelerators. The multiplication-free design enables higher operating frequencies and improved frame rates while maintaining architectural simplicity.

5. Conclusion

We present BDLUT, integrating blind denoising with hardware-optimized LUTs for efficient image restoration. Our approach achieves SOTA performance among LUT methods while maintaining minimal computational overhead. The successful FPGA implementation demonstrates BDLUT's practicality for real-world applications, offering an effective solution for resource-constrained edge computing scenarios.

Acknowledgments

This work is supported in part by the HKU-TCL Joint Research Centre for Artificial Intelligence, and in part by the General Research Fund (GRF) project 17203224 and the Theme-based Research Scheme (TRS) project T45-701/22-R of the Research Grants Council (RGC), Hong Kong SAR.

References

[1] Fanny Spagnolo, Pasquale Corsonello, Fabio Frustaci, and Stefania Perri. 2023. Design of Approximate Bilateral Filters for Image Denoising on FPGAs. *IEEE Access* 11 (2023), 1990–2000. https://doi.org/10.1109/ACCESS.2022.3233921

[2] Fanny Spagnolo, Pasquale Corsonello, Fabio Frustaci, and Stefania Perri. 2024. KIT: Kernel Isotropic Transformation of Bilateral Filters for Image Denoising on FPGA. In *2024 34th International Conference on Field-Programmable Logic and Applications (FPL)*. 19–23. https://doi.org/10.1109/FPL64840.2024.00013

[3] Ruoheng Yao, Lei Chen, Pingcheng Dong, Zhuoyu Chen, and Fengwei An. 2022. A Compact Hardware Architecture for Bilateral Filter With the Combination of Approximate Computing and Look-Up Table. *IEEE Transactions on Circuits and Systems II: Express Briefs* 69, 7 (2022), 3324–3328. https://doi.org/10.1109/TCSII.2022.3159261

[4] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. 2007. Color Image Denoising via Sparse 3D Collaborative Filtering with Grouping Constraint in Luminance-Chrominance Space. In *2007 IEEE International Conference on Image Processing*, Vol. 1. I – 313–I – 316. https://doi.org/10.1109/ICIP.2007.4378954

[5] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. 2017. Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising. *IEEE Transactions on Image Processing* 26, 7 (2017), 3142–3155. https://doi.org/10.1109/TIP.2017.2662206

[6] Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. 2021. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*. 1833–1844.

[7] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. 2021. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF international conference on computer vision*. 1905–1914.

[8] Younghyun Jo and Seon Joo Kim. 2021. Practical Single-Image Super-Resolution Using Look-Up Table. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 691–700. https://doi.org/10.1109/CVPR46437.2021.00075

[9] Jiacheng Li, Chang Chen, Zhen Cheng, and Zhiwei Xiong. 2022. MuLUT: Cooperating Multiple Look-Up Tables for Efficient Image Super-Resolution. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVIII*. 238–256. https://doi.org/10.1007/978-3-031-19797-0_14

[10] Guandu Liu, Yukang Ding, Mading Li, Ming Sun, Xing Wen, and Bin Wang. 2023. Reconstructed Convolution Module Based Look-Up Tables for Efficient Image Super-Resolution . In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. 12183–12192. https://doi.org/10.1109/ICCV51070.2023.01122

[11] Yinglong Li, Jiacheng Li, and Zhiwei Xiong. 2024. Look-Up Table Compression for Efficient Image Restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 26016–26025.

[12] Binxiao Huang, Jason Chun Lok Li, Jie Ran, Boyu Li, Jiajun Zhou, Dahai Yu, and Ngai Wong. 2024. Hundred-Kilobyte Lookup Tables for Efficient Single-Image Super-Resolution. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, Kate Larson (Ed.). International Joint Conferences on Artificial Intelligence Organization, 857–865. https://doi.org/10.24963/ijcai.2024/95 Main Track.

[13] Xuyang Duan, Ruiqi Xie, and Jun Han. 2021. An Energy-Efficient Image Denoising Accelerator with Depth-wise Separable Convolution and Fused-Layer Architecture. In *2021 IEEE 14th International Conference on ASIC (ASICON)*. 1–4. https://doi.org/10.1109/ASICON52560.2021.9620485

[14] Jun Xu, Lei Zhang, David Zhang, and Xiangchu Feng. 2017. Multi-Channel Weighted Nuclear Norm Minimization for Real Color Image Denoising. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

[15] Gu Shuhang, Zhang Lei, Zuo Wangmeng, and Feng Xiangchu. 2014. Weighted Nuclear Norm Minimization with Application to Image Denoising. In *Proceedings of the IEEE conference on computer vision and pattern recognitionn (CVPR)*.

[16] Eirikur Agustsson and Radu Timofte. 2017. NTIRE 2017 Challenge on Single Image Super-Resolution: Dataset and Study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.

[17] D. Martin, C. Fowlkes, D. Tal, and J. Malik. 2001. A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics. In *Proc. 8th Int'l Conf. Computer Vision*, Vol. 2. 416–423.

[18] Rich Franzen. 1999. Kodak Lossless True Color Image Suite. Retrieved Dec 18, 2024 from https://r0k.us/graphics/kodak/

[19] Zhang Lei, Wu Xiaolin, Buades Antoni, and Li Xin. 2011. Color Demosaicking by Local Directional Interpolation and Non-local Adaptive Thresholding. *Journal of Electronic Imaging* 20, 2 (2011), 023016. https://doi.org/10.1117/1.3600632

[20] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. 2015. Single Image Super-Resolution From Transformed Self-Exemplars. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5197–5206.

[21] Abdelhamed Abdelrahman, Lin Stephen, and Brown Michael S. 2018. A High-Quality Denoising Dataset for Smartphone Cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*.