arXiv:2412.12942v1 [eess.IV] 17 Dec 2024

Transforming Single Photon Camera Images to Color High Dynamic Range Images

Sumit Sharma[1 [0009−0004−7454−3793]], Girish Rongali[2 [0009−0005−2239−6821]], and Kaushik Mitra[3 [0000−0001−6747−9050]]

Indian Institute of Technology, Madras,India

Abstract. Traditional CMOS sensors suffer from restricted dynamic range and sub-optimal performance under extreme lighting conditions. They are affected by electronic noise in low light conditions and pixel saturation while capturing high illumination. Recent High Dynamic Range (HDR) imaging methods, often designed for CMOS sensors, attempt to address these issues by fusing multiple exposures. However, they frequently introduce artifacts like ghosting and light flickering in dynamic scenarios and non-uniform signal-to-noise ratios in extreme dynamic range conditions. Recently, Single Photon Avalanche Diode (SPAD) sensors, also known as Single Photon Camera (SPC) sensors, have surpassed CMOS sensors due to their capability to capture individual photons with high timing precision. Unlike traditional digital cameras that first convert light energy into analog electrical currents and then digitize them, SPAD sensors perform direct photon detection, making them less susceptible to extreme illumination conditions. Their distinctive non-linear response curve aids in capturing photons across both low-light and high-illumination environments, making them particularly effective for high-dynamic range imaging. Despite their advantages, images from SPAD sensors are often noisy and visually unappealing. To address these challenges, our paper evaluates state-of-the-art architectures for converting monochromatic SPAD images into color HDR images at various resolutions. We present a two-stage approach: first, transforming monochromatic SPAD images into color images via Image2Image translation networks, and second, enhancing these color images to achieve HDR via Single Image HDR methods. Our evaluation involves both qualitative and quantitative assessments of these architectures, focusing on their effectiveness in each stage of the conversion process.

Keywords: Single Photon Avalanche Diode · High Dynamic Range Imaging· Generative Adversarial Networks.

1 Introduction

Traditional imaging systems often struggle with scenes featuring high levels of illumination and intricate details, such as bright light sources and highlights. Conventional image sensors have limited dynamic range because of limited full-well capacity. This leads to saturation at high brightness levels. To overcome

the constraints of traditional sensors, several computational and hardware approaches have been developed [1] for capturing High Dynamic Range (HDR) images. One common method is exposure bracketing, which combines multiple images taken at different exposure levels into a single HDR image [2],[3]. However, this method can introduce ghosting artifacts [4] and light flicker in dynamic scenes [7], [6], [5]. To mitigate these issues, commercial HDR sensors typically use a limited number of exposures—usually between two and four—either sequentially or through dual-pixel architectures [6],[12]. Despite these improvements, HDR cameras using CMOS sensors still face challenges such as motion artifacts and a reduced signal-to-noise ratio (SNR) in scenes with extreme dynamic ranges. The process of recovering an HDR image from a few exposures can lead to uneven SNR artifacts, making it difficult to preserve fine details and effective denoising of the image [14], [13], [7].

In contrast, Single Photon avalanche Diode (SPAD) sensors offer a promising solution by measuring brightness through precise timing of individual photons based on Poisson statistics. This approach allows SPAD sensors to surpass the saturation limits of traditional pixels and capture timing information with high accuracy [8]. Recent advancements in Single-Photon Cameras (SPC) have demonstrated their superior dynamic range compared to conventional CMOS sensors [9]. SPAD sensors excel in both low and high light conditions within a single exposure, enabling them to produce detailed images across a wide range of brightness levels [10], [8]. Despite these advantages, SPC images are not visually appealing, are noisy, and have limited resolutions. We propose two-stage as well as single stage approaches for converting monochromatic SPC images captured by single-photon cameras into high-resolution color HDR images. In the two-stage approach, we first utilize existing Image-to-Image translation networks to convert monochromatic SPC images to color LDR images and then convert them to color HDR images based on single image HDR reconstruction techniques. We evaluate various two-stage architectures through comprehensive quantitative and qualitative analyses. We also evaluate a direct image2image translation method for converting monochromatic SPC images to color HDR images. Our main contributions are:

- Transforming SPC images to color HDR images.
- Detailed analysis of various two-stage architectures for achieving the above goal.
- Single-stage Image2Image translation from monochromatic SPC images to color HDR images at high resolution produces comparable results.

2 Related Work

2.1 Image-to-Image Translation

Given that Single Photon Camera (SPC) images lack color information, employing image-to-image translation methods presents a viable avenue for colorizing SPC images. Among these methods, Generative Adversarial Networks (GANs)

stand out as a prominent framework for estimating generative models through an adversarial process. In this process, two models are concurrently trained: a Generative model 'G' that captures the data distribution, and a Discriminative model 'D' that estimates the probability that a sample originates from the training data rather than 'G' [15]. Building upon this concept [15], researchers have delved into conditional adversarial networks as a versatile solution to various image-to-image translationI2I challenges. These networks not only learn the mapping from input images to output images but also acquire a loss function to facilitate this mapping training [16]. Defining and optimizing perceptual loss functions based on high-level features extracted from pre-trained networks makes it feasible to generate high-quality images [17]. Consequently, the fusion of training feed-forward convolutional neural networks with per-pixel loss and perceptual loss functions has emerged as a favored approach for numerous image-to-image translation tasks [18],[19],[20]. Furthermore, various unsupervised methods have been proposed for tackling the image-to-image translation task [21],[22],[23,24].In contrast, one approach [25] focuses on producing images with a photographic appearance that aligns with the input layout. This approach operates as a rendering engine, translating a two-dimensional semantic specification of the scene into a photographic image without relying on adversarial training, which may suffer from training instability and optimization issues. To circumvent these challenges, this approach [25] adopts a direct regression objective based on computing distances between image features extracted by deep neural networks [26], [27], [17]. Consequently, although this approach [25] yields the first model capable of synthesizing 2048 x 1024 images, it may lack finer details and textures. Alternatively, another approach [28] achieves visually appealing results at a resolution of 2048x1024 by incorporating a novel adversarial loss function, along with multi-scale generator and discriminator architectures. Unpaired image-to-image translation often relies on maximizing mutual information between source and translated images across different domains to preserve source content and prevent unnecessary modifications. Self-supervised contrastive learning has proven effective in this context, where it constrains features to ensure fidelity to the source. However, prior methods often use features from random locations to impose constraints, potentially lacking sufficient information from the source domain.

2.2 Single Image HDR

As a result of the non-linear response curve inherent in the SPAD pixel, characterized by an asymptotic saturation limit, and its efficacy in operating under exceptionally high flux levels, Single Photon Counting (SPC) devices exhibit the capability to capture High Dynamic Range (HDR) images even in challenging lighting conditions, such as those encountered in extremely low-light and high-light scenarios. Consequently, incorporating SPAD technology within single-image HDR methodologies is a promising avenue for achieving vibrant and color-rich HDR imagery [10],[9]. Reconstruction of high dynamic range (HDR) images through the fusion of multiple exposures stands as a conventional technique, harnessing the distinctive information encapsulated within

each exposure[2],[29].Specifically, low exposures encapsulate regions abundant in intensity, whereas high exposures capture details within low-intensity areas, a practice commonly referred to as Exposure Bracketing. Nonetheless, challenges arise within this methodology when tackled with motion scenarios, often resulting in artifacts such as ghosting. In such circumstances, Single-image HDR techniques present a promising alternative that is particularly adept at addressing issues like misalignment. Eilertsen et al. [30] introduced a novel approach for reconstructing HDR images from low dynamic range (LDR) input images by inferring missing information within luminance-rich segments, such as highlights, which are lost due to limitations of the camera sensor. This was achieved through a fully convolutional neural network (CNN) structured in the form of a hybrid dynamic range autoencoder. Similarly, Yang et al. [31] conceptualized the image correction task as an HDR transformation process, initially reconstructing absent details within the HDR domain, followed by tone-mapping to generate the output LDR image with reinstated details. In contrast, Moriwaki et al. [32] devised a method aimed at minimizing a hybrid loss, comprising perceptual and adversarial losses in addition to HDR reconstruction loss. Liu et al. [33] modeled the HDR-to-LDR image formation pipeline, describing dynamic range clipping, non-linear mapping from a camera response function, and quantization. They further employed three specialized CNNs to reverse these steps, imposing effective physical constraints to facilitate the training of individual sub-networks. On the other hand, Santo et al. [34] reconstructed HDR images by recovering saturated pixels from an input LDR image through a feature masking mechanism, thereby reducing the influence of features from saturated areas. To synthesize visually pleasing textures, they employed a VGG-based perceptual loss function. The light-weight Deep Neural Network [42] has also been specifically designed to address noise and compression in Deep Neural Network-based Single Image High Dynamic Range (SIHDR) processing of legacy Standard Dynamic Range Images.

3 Proposed work

3.1 Single Photon Camera Soft Saturation

Conventional sensor pixels typically exhibit a linear input-output response, leading to hard saturation when subjected to a sudden increase in incident flux. In contrast, Single Photon Camera (SPC) pixels demonstrate a non-linear response characterized by an asymptotic saturation threshold. Upon detecting each photon, the Single Photon Avalanche Diode (SPAD) enters a predefined dead-time interval, rendering it temporarily insensitive to subsequent photons. This inherent property introduces non-linearity into the response curve of the SPAD, causing it to miss a fraction of incident photons as flux escalates in its passive mode. The theoretical saturation limit of an SPC pixel in passive mode significantly surpasses that of conventional sensors, enabling it to resist saturation even under extreme brightness conditions. Instead of hard saturation, the SPC pixel reaches a soft saturation point where it remains responsive, albeit with a

diminished signal-to-noise ratio. This soft saturation threshold notably exceeds the saturation limits observed in conventional sensors, making SPC pixels particularly well-suited for applications demanding extremely High Dynamic Range Imaging capabilities.

3.2 Single Photon Camera Simulator

Due to the continuous advancement of research in SPC and their functionalities, their availability in the commercial market remains limited. Hence, we prepare a Single Photon Camera simulator leveraging the insights from recent studies [36],[9]. This simulator is designed to replicate SPC image generation by employing ground-truth HDR images[37] as input data. Essentially, it replicates real-life scene dynamics as if captured by a single-photon camera. The SPC Simulator has been developed under the assumption of the SPC operating in passive mode. With each photon detection event, the SPAD undergoes a dead time interval, during which it remains inactive and does not capture any photons. It is assumed that the arrival of each photon incident per pixel follows Poisson statistics. Meanwhile, the number of detected photons 'N_{SPC}^{T}' within a fixed exposure time 'T' follows a renewal process, as detailed in Grimmett and Stirzaker's work on probability [38]. Consequently, the mean and variance are determined by the methodology outlined in Ingle and Proakis's publication [10] as:

$$E[N_T^{SPC}] = q_{SPAD}\phi T/[1 + q_{SPAD}\phi\tau_d] \tag{1}$$

$$Var[N_T^{SPC}] = q_{SPAD}\phi T/[1 + q_{SPAD}\phi\tau_d]^3, \tag{2}$$

where ϕ is the photon flux, that is, number of photons per sec, q_{SPAD} is the quantum efficiency of the SPC camera. Finally, we also add an averaging step in order to generate different numbers of averaged frames[36] in the '.hdr' format. The simulation begins by loading the High Dynamic Range (HDR) image using the OpenCV library. Once the HDR image is loaded, it is converted to gray-scale to facilitate the estimation of the luminance present within the image. This luminance is subsequently utilized to determine the appropriate exposure time, denoted as τ_d, for HDR images to ensure optimal visibility. The gray-scale-converted image is treated as a **photon flux** 'ϕ', which is then used to calculate the **mean photon count** and **variance photon count** in accordance with the equations 1 and 2 outlined in [9].

During simulation, adjustments are made to key parameters such as exposure time (T) and SPC pixel sensitivity (q_{SPAD}) while maintaining a fixed dead time. The objective is to ensure visibility of all regions simulated from the SPC simulator, which would facilitate the attainment of an optimal reconstructed HDR image [9]. In order to get the optimal results, we select the exposure time based on the luminance of the ground truth HDR images. Different resolution of simulator outputs can be seen from the 1. We save the monochromatic SPC images obtained from the SPC simulator as 8-bit monochromatic images. This process was essential to prepare the input monochromatic SPC images for training and testing.

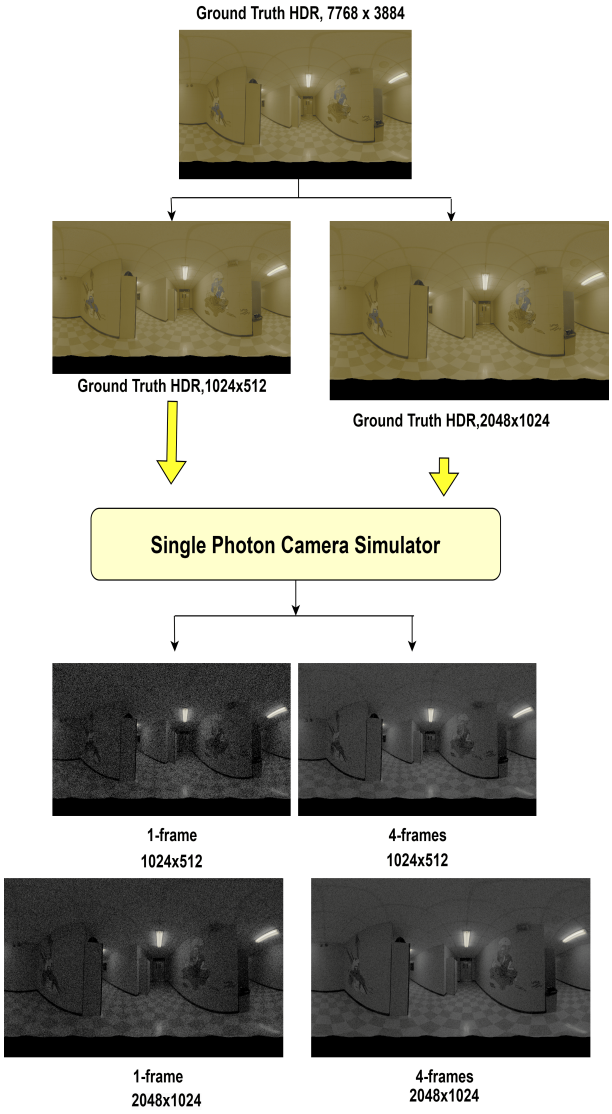


Fig. 1: SPC simulation pipeline: Different averaged SPC frames at different resolutions are generated from the SPC simulator based on luminance-based exposure time selection.

3.3 Two-stage and Single-stage approaches

SPC have limitations in resolution and color capabilities. However, Canon has developed a small SPAD sensor (13.2 mm x 9.9 mm) that can capture 3.2 megapixels of color images [41]. While there are potential hardware solutions, current



Fig. 2: This study examines various methods for converting a monochromatic SPC image to a Color HDR image using state-of-the-art techniques. We explore Pix2Pix, Pix2PixHD, and CycleGAN for the translation from monochromatic to color. Subsequently, we investigate DeepHDR, LHDR, SIHDR, and ExpandNet for converting the colorized image to its HDR version.

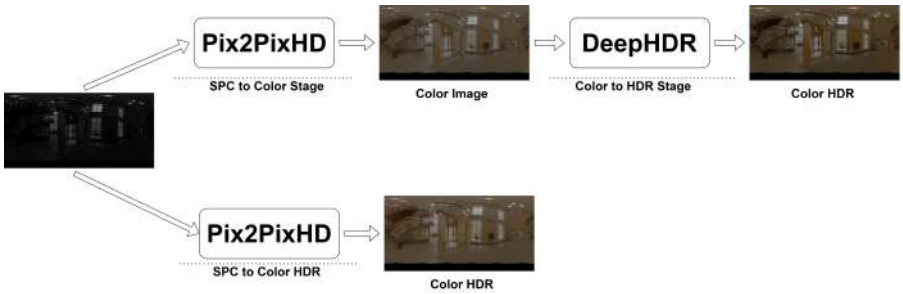


Fig. 3: Best possible combination among the different state-of-the-art architectures. For image2image translation, we found Pix2PixHD performs best at low and high resolutions, and DeepHDR performs best when creating HDR versions among different state-of-the-art architectures. We also modified Pix2PixHD, for direct translation i.e., moving from Monochromatic HDR to Color HDR and getting Comparable results

prototypes are not ready to compete directly with CMOS image sensors used in smartphones. These sensors have strict requirements for pixel size, power use, and processing efficiency, which current SPAD technology does not yet meet [40].

This paper explores different computational imaging techniques to convert SPC monochromatic images into color high dynamic range (HDR) images. We explore two approaches for converting monochromatic SPC images to color HDR images: i) two-stage models and ii) single-stage models. In two-stage models, we explore different combinations of image-to-image SPC colorization methods and single-image HDR reconstruction methods. This is shown in Figure 2. We explore Pix2PixHD[28],Pix2Pix[16] and CycleGAN[46] for colorization task and Expandnet[45], DeepHDR[34], SIHDR[33] and LHDR[42], for single image HDR reconstruction. From Figure 2, it is clear that a total of 12 combinations is possible to move from a SPC monochromatic image to a colorful HDR version.

To enhance the clarity of details in monochromatic SPC images such as text, we extended our study by applying averaging techniques. Specifically, we experimented with averaging different frames ranging from 1 to a maximum of 4. We

create a paired dataset consisting of SPC simulated images, which are different frames averaged monochromatic SPC images and corresponding logarithmic tonemapped color images of different resolutions, i.e., 1024x512 and 2048x1024 pixels. We train the different image2image translation networks on these paired datasets for monochromatic SPC to color conversion. For the subsequent stage of color to HDR conversion, we used pre-trained models of various Single Image HDR methods.

Finally, we also tried single-stage translation to move from monochromatic image to color HDR image. For this, we trained Pix2PixHD[28] between monochromatic SPC image and 32-bits-per-channel ground truth HDR images. The synthesized images are normalized images in the range 0-1, which are inverse tonemapped images with gamma correction. We then post-process those images to 32-bits floating point images. These images were directly compared with 32-bits-per-channel ground truth HDR images in terms of HDR VDP scores.

4 Experimentation

All the experiments are done on the Laval indoor dataset [47] for our analysis, which contains a total of 2362 high-resolution HDR images of resolution 7768x3884. The images are downsampled to a resolution of 1024x512 and 2048x1024 to reduce the computational complexity. We selected these two different resolutions to assess performance improvements relative to image resolution. Monochromatic SPC frames, simulated from these downsampled images, serve as the input to the model. Since the SPC frames are very noisy, we also averaged a few frames (for example 4 frames) and gave them as input to the model, which we referred to as 4-frame averaged SPC images. As image translation methods cannot predict high bit-depth images directly, the monochromatic SPC images are tone-mapped to produce 8-bit monochromatic images, which are the input to the two-stage models. For training the colorization models in the two-stage models, we tonemapped the original color HDR images to produce color LDR images, which are then used as ground-truth for training the colorization networks. After extensive experimentation with Reinhard [1] and Adaptive logarithmic tone mapping, we selected logarithmic tone mapping due to its superior performance. For the second stage of the two-stage approach, we use different pre-trained state-of-the-art single-image HDR reconstruction models to convert the colorized SPC images to color HDR images. All the experimentations are done using two NVIDIA GeForce RTX 3090 GPUs.

5 Quantitative and Qualitative Analysis

We perform quantitative analysis for both the stages of the two-stage apporaches: i) monochromatic SPC to colorized SPC and ii) colorized SPC to color HDR. For the task of SPC colorization, we used *psnr*,*ssim*, and *lpips* as metrics. For the task of converting color SPC to color HDR images, we use *HDR-VDP3* [43] metric.

Table 1 shows the comparison of various image-to-image translation models such as Pix2PixHD[28], Pix2Pix[16], and CycleGAN[46] for the task of SPC colorization. The table illustrates the results with single frame and 4-frame (averaged) for 1024x512 and 2048x1024 resolutions. From the table 1, it is clear that Pix2PixHD [28] achieves the highest performance for the colorization task among state-of-the-art methods, as measured by LPIPS, PSNR, and SSIM scores. Averaging 4 frames yields better results than using a single SPC monochromatic frame.

For the task of generating HDR color images, we utilized single-image HDR reconstruction methods, including DeepHDR[34], SIHDR[33], ExpandNet[45], and LHDR[42]. From Table 2, it is clear that the combination of Pix2PixHD[28] and DeepHDR[34] is quantitatively giving the best possible HDR reconstruction with the highest HDR-VDP3 scores among all methods. Also, the performance of these methods is better at higher resolutions of 2048x1024.

Qualitatively, by observing the reconstructed images from different image-to-image translation networks in Figure 5, it is clear that Pix2PixHD[28] produces better results than other state-of-the-art architectures. We can also observe from figure 4 that CycleGAN[46] fails to preserve colors and also fails to reconstruct the light sources faithfully at higher resolutions.

We also experimented with Pix2PixHD [28] for the task of single-stage translation from Monochromatic SPC images to Color HDR images. Given that Pix2PixHD produced superior results among all models for the two-stage translation task, we modified its architecture as described in 3.3 for that task of single-stage translation. As shown in Table 3, the results obtained using a single-stage translation task are slightly comparable but less than those obtained using two-stage translation. Figure 6 illustrates that the colors in the HDR image produced by the two-stage translation (Pix2PixHD with DeepHDR) are closely aligned with the ground truth, whereas the colors generated by the modified Pix2PixHD in the single-stage translation deviate slightly from the ground truth.Please refer to Figure 3 for a comparison of direct and indirect approaches, illustrating the optimal combination of these methods. Single Image HDR techniques are primarily designed to convert Low Dynamic Range (LDR) images to High Dynamic Range (HDR) images, which typically involves expanding the bit depth of each color channel. However, for the task of translating monochromatic Single-Photon Camera (SPC) images to color HDR images, additional losses are required to effectively predict the color information. In this context, **adversarial loss** can be a promising approach [28].

6 Conclusion and future works

This paper presents a detailed analysis of transforming single-photon camera (SPC) captured images into colorful HDR images using two-stage architecture. SPC images inherently possess high dynamic range due to their soft-saturation characteristics [9], [8], [10]. Our simulator [9] is capable of generating single-frame monochromatic HDR images by adjusting parameters such as SPC sensitivity,

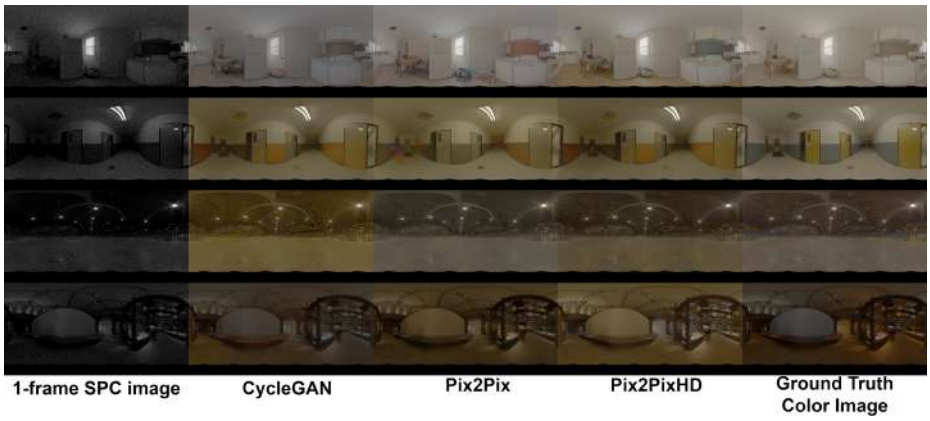


Fig. 4: This figure presents a qualitative analysis of 1-frame SPC to color translation at a resolution of 2048x1024. It displays various colorized outputs from image-to-image translation methods, with the rightmost column showing the logarithmic tone-mapped ground truth color images. Notably, Pix2PixHD yields the most favorable results. Please zoom in for further detail.



Fig. 5: This figure presents qualitative analysis of four-frame averaged SPC for color translation at a resolution of 2048x1024. It displays various image-to-image translation outputs. The rightmost column features logarithmic tone-mapped ground truth color images. Notably, the best results are achieved using Pix2PixHD.Please zoom in for further details

Table 1: SPC monochromatic to colorization results. Pix2PiHD gives the best result.

Method	Resolution	lpips 1-frame↓	lpips 4-frames↓	PSNR 1-frame↑	PSNR 4-frames↑	SSIM 1-frame↑	SSIM 4-frames↑
Pix2PixHD[28]	1024x512	**0.2417**	**0.2029**	**25.34**	**25.36**	**0.8456**	**0.8718**
Pix2Pix[16]	1024x512	0.4102	0.3901	24.87	25.09	0.8279	0.848
CycleGAN[46]	1024x512	0.4068	0.3973	23.1	23.25	0.8192	0.8235
Pix2PixHD[28]	2048x1024	**0.2254**	**0.1995**	**26.53**	**26.5**	**0.887**	**0.9056**
Pix2Pix[16]	2048x1024	0.4418	0.443	25.25	25.23	0.8641	0.8701
CycleGAN [46]	2048x1024	0.4498	0.4547	23.2	23.07	0.8384	0.8364

Table 2: SPC monochromatic to HDR results. The HDR VDP scores are averaged scores over 262 images. Pix2PixHD + DeepHDR gives the best result.

Method	Resolutions	Q score (1-frame) ↑		Q score (4-frames) ↑	
		1024x512	2048x1024	1024x512	2048x1024
Pix2PixHD[28] + Expandnet [45]	Low+High	3.1596	3.41	3.2667	3.4533
Pix2PixHD[28] + DeepHDR[34]	Low+High	**3.6388**	**3.9078**	**3.7393**	**4.0079**
Pix2PixHD[28] + SIHDR [33]	Low+High	2.7251	3.1357	2.824	3.2111
Pix2PixHD[28] + LHDR [42]	Low+High	3.1131	3.386	3.2056	3.4339
Pix2Pix[16] + Expandnet [45]	Low+High	3.0453	3.1783	3.1666	3.2977
Pix2Pix[16] + DeepHDR [34]	Low+High	3.4545	3.5753	3.5624	3.6699
Pix2Pix[16] + SIHDR [33]	Low+High	2.8896	3.0503	3.0131	3.1685
Pix2Pix[16] + LHDR [42]	Low+High	3.1354	3.2717	3.2344	3.3639
CycleGAN[46] + Expandnet [45]	Low+High	2.9679	3.1897	3.1041	3.2136
CycleGAN[46] + DeepHDR [34]	Low+High	3.3189	3.5239	3.4861	3.581
CycleGAN[46] + SIHDR [33]	Low+High	2.7932	3.0345	2.9699	3.1025
CycleGAN[46] + LHDR [42]	Low+High	3.0445	3.2358	3.118	3.2432

Table 3: Single stage translation from SPC monochromatic to color HDR

Method	Resolution	1-frame,Q score	4-frames,Q score
Pix2PixHD	1024x512	3.4868	3.6154
Pix2PixHD	2048x1024	3.8355	3.9734

exposure time, and a **fixed dead time** of **150 ns**. The exposure time is determined based on the **luminance** present in the image. Our two-stage approach involves image-to-image translation methods for colorization and single-image HDR reconstruction methods for generating the final HDR version. We have also compared the two-stage approach with a single-stage approach and obtained comparable results.

Given that single photon cameras can record high frame rates exceeding 96k frames per second, this research can be extended to develop high-frame-rate HDR videos, which will enhance our understanding of various scenes.

7 Acknowledgement

This work was supported in part by IITM Pravartak Technologies Foundation.

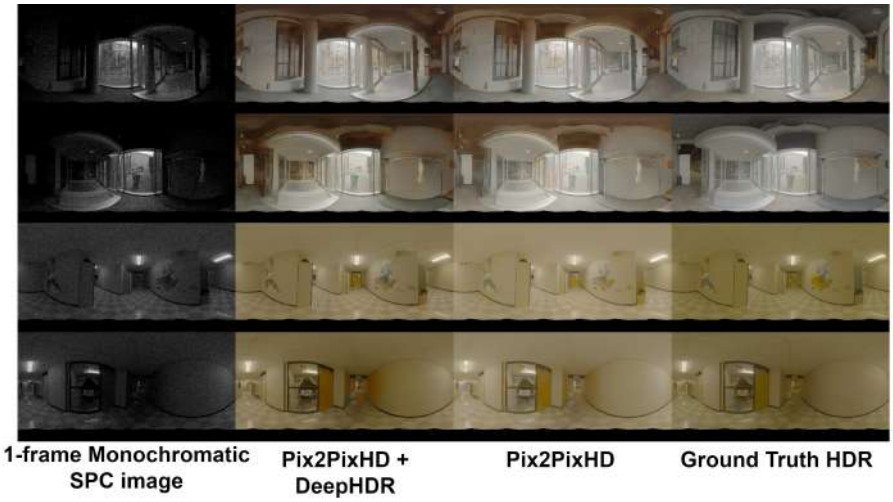


Fig. 6: Qualitative analysis of 1-frame SPC for color HDR translation at a resolution of 2048x1024 is presented. This figure displays outputs from both the two-stage (Pix2PixHD + DeepHDR) and single-stage (Pix2PixHD) methods. The rightmost column shows the logarithmic tone-mapped ground truth HDR images. Results demonstrate comparability between the approaches.Please zoom in for further details

References

1. Reinhard, Erik. "High dynamic range imaging." In Computer Vision: A Reference Guide, pp. 1-6. Cham: Springer International Publishing, 2020.
2. Debevec, Paul E., and Jitendra Malik. "Recovering high dynamic range radiance maps from photographs." In Seminal Graphics Papers: Pushing the Boundaries, Volume 2, pp. 643-652. 2023.
3. Mann, S., and R. Picard. "Beingundigital'with digital cameras." MIT Media Lab Perceptual 1, no. 2 (1994): 3.
4. Tursun, Okan Tarhan, Ahmet Oğuz Akyüz, Aykut Erdem, and Erkut Erdem. "The state of the art in HDR deghosting: A survey and evaluation." In Computer Graphics Forum, vol. 34, no. 2, pp. 683-707. 2015.
5. Deegan, Brian. "LED flicker: Root cause, impact and measurement for automotive imaging applications." Electronic Imaging 30 (2018): 1-6.
6. Iida, S., Y. Sakano, T. Asatsuma, M. Takami, I. Yoshiba, N. Ohba, H. Mizuno et al. "A 0.68 e-rms random-noise 121dB dynamic-range sub-pixel architecture CMOS image sensor with LED flicker mitigation." In 2018 IEEE International Electron Devices Meeting (IEDM), pp. 10-2. IEEE, 2018.
7. Velichko, Sergey, Scott Johnson, Dan Pates, Chris Silsby, Cornelis Hoekstra, Ray Mentzer, and Jeff Beck. "140 db dynamic range sub-electron noise floor image sensor." Proceedings of the IISW 1 (2017).
8. Ingle, Atul, Trevor Seets, Mauro Buttafava, Shantanu Gupta, Alberto Tosi, Mohit Gupta, and Andreas Velten. "Passive inter-photon imaging." In Proceedings of the

IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8585-8595. 2021.

9. Liu, Yuhao, Felipe Gutierrez-Barragan, Atul Ingle, Mohit Gupta, and Andreas Velten. "Single-photon camera guided extreme dynamic range imaging." In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 1575-1585. 2022.
10. Ingle, Atul, Andreas Velten, and Mohit Gupta. "High flux passive imaging with single-photon sensors." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 6760-6769. 2019.
11. Fattal, Raanan, Dani Lischinski, and Michael Werman. "Gradient domain high dynamic range compression." In Seminal Graphics Papers: Pushing the Boundaries, Volume 2, pp. 671-678. 2023.
12. Asatsuma, T., Y. Sakano, S. Iida, M. Takami, I. Yoshiba, N. Ohba, H. Mizuno et al. "Sub-pixel architecture of cmos image sensor achieving over 120 db dynamic range with less motion artifact characteristics." In Proceedings of the 2019 International Image Sensor Workshop, vol. 1. 2019.
13. Mase, Mitsuhito, Shoji Kawahito, Masaaki Sasaki, Yasuo Wakamori, and Masanori Furuta. "A wide dynamic range CMOS image sensor with multiple exposure-time signal outputs and 12-bit column-parallel cyclic A/D converters." IEEE Journal of Solid-State Circuits 40, no. 12 (2005): 2787-2795.
14. Yang, David XD, and Abbas El Gamal. "Comparative analysis of SNR for image sensors with enhanced dynamic range." In Sensors, cameras, and systems for scientific/industrial applications, vol. 3649, pp. 197-211. SPIE, 1999.
15. Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. "Generative adversarial nets." Advances in neural information processing systems 27 (2014).
16. Isola, Phillip, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. "Image-to-image translation with conditional adversarial networks." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1125-1134. 2017.
17. Johnson, Justin, Alexandre Alahi, and Li Fei-Fei. "Perceptual losses for real-time style transfer and super-resolution." In Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II 14, pp. 694-711. Springer International Publishing, 2016.
18. Dong, Hao, Simiao Yu, Chao Wu, and Yike Guo. "Semantic image synthesis via adversarial learning." In Proceedings of the IEEE international conference on computer vision, pp. 5706-5714. 2017.
19. Kaneko, Takuhiro, Kaoru Hiramatsu, and Kunio Kashino. "Generative attribute controller with conditional filtered generative adversarial networks." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 6089-6098. 2017.
20. Karacan, Levent, Zeynep Akata, Aykut Erdem, and Erkut Erdem. "Learning to generate images of outdoor scenes from attributes and semantic layouts." arXiv preprint arXiv:1612.00215 (2016).
21. Bousmalis, Konstantinos, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. "Unsupervised pixel-level domain adaptation with generative adversarial networks." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3722-3731. 2017.
22. Liu, Ming-Yu, Thomas Breuel, and Jan Kautz. "Unsupervised image-to-image translation networks." Advances in neural information processing systems 30 (2017).

23. Shrivastava, Ashish, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. "Learning from simulated and unsupervised images through adversarial training." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2107-2116. 2017.
24. Taigman, Yaniv, Adam Polyak, and Lior Wolf. "Unsupervised cross-domain image generation." arXiv preprint arXiv:1611.02200 (2016).
25. Chen, Qifeng, and Vladlen Koltun. "Photographic image synthesis with cascaded refinement networks." In Proceedings of the IEEE international conference on computer vision, pp. 1511-1520. 2017.
26. Dosovitskiy, Alexey, and Thomas Brox. "Generating images with perceptual similarity metrics based on deep networks." Advances in neural information processing systems 29 (2016).
27. Gatys, Leon A., Alexander S. Ecker, and Matthias Bethge. "Image style transfer using convolutional neural networks." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2414-2423. 2016.
28. Wang, Ting-Chun, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. "High-resolution image synthesis and semantic manipulation with conditional gans." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 8798-8807. 2018.
29. Mertens, Tom, Jan Kautz, and Frank Van Reeth. "Exposure fusion: A simple and practical alternative to high dynamic range photography." In Computer graphics forum, vol. 28, no. 1, pp. 161-171. Oxford, UK: Blackwell Publishing Ltd, 2009.
30. Gabriel, Eilertsen. "HDR image reconstruction from a single exposure using deep CNNs." ACM Trans. Graph. 36 (2017): 178-1.
31. Yang, Xin, Ke Xu, Yibing Song, Qiang Zhang, Xiaopeng Wei, and Rynson WH Lau. "Image correction via deep reciprocating HDR transformation." In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1798-1807. 2018.
32. Moriwaki, Kenta, Ryota Yoshihashi, Rei Kawakami, Shaodi You, and Takeshi Naemura. "Hybrid loss for learning single-image-based HDR reconstruction." arXiv preprint arXiv:1812.07134 (2018).
33. Liu, Yu-Lun, Wei-Sheng Lai, Yu-Sheng Chen, Yi-Lung Kao, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. "Single-image HDR reconstruction by learning to reverse the camera pipeline." In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 1651-1660. 2020.
34. Santos, Marcel Santana, Tsang Ing Ren, and Nima Khademi Kalantari. "Single image HDR reconstruction using a CNN with masked features and perceptual loss." arXiv preprint arXiv:2005.07335 (2020).
35. Le, Phuoc-Hieu, Quynh Le, Rang Nguyen, and Binh-Son Hua. "Single-image hdr reconstruction by multi-exposure generation." In Proceedings of the IEEE/CVF winter conference on applications of computer vision, pp. 4063-4072. 2023.
36. Goyal, Bhavya, and Mohit Gupta. "Photon-starved scene inference using single photon cameras." In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 2512-2521. 2021.
37. Bolduc, Christophe, Justine Giroux, Marc Hébert, Claude Demers, and Jean-François Lalonde. "Beyond the pixel: a photometrically calibrated hdr dataset for luminance and color prediction." In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 8071-8081. 2023.
38. Grimmett, Geoffrey, and David Stirzaker. Probability and random processes. Oxford university press, 2020.

39. Hu, Xueqi, Xinyue Zhou, Qiusheng Huang, Zhengyi Shi, Li Sun, and Qingli Li. "Qs-attn: Query-selected attention for contrastive learning in i2i translation." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 18291-18300. 2022.
40. Ma, Sizhuo, Varun Sundar, Paul Mos, Claudio Bruschini, Edoardo Charbon, and Mohit Gupta. "Seeing Photons in Color." ACM Transactions on Graphics (TOG) 42, no. 4 (2023): 1-16.
41. "Canon Has Successfully Developed an Advanced SPAD Sensor." Canon Global. Last modified 2023. https://global.canon/en/technology/spad-sensor-2023.
42. Guo, Cheng, and Xiuhua Jiang. "Lhdr: Hdr reconstruction for legacy content using a lightweight dnn." In Proceedings of the Asian Conference on Computer Vision, pp. 3155-3171. 2022.
43. Mantiuk, Rafal K., Dounia Hammou, and Param Hanji. "HDR-VDP-3: A multi-metric for predicting image differences, quality and contrast distortions in high dynamic range and regular content." arXiv preprint arXiv:2304.13625 (2023).
44. Hatchett, Dr. Jonathan, Fatih Bakir, Göksu Güvendiren, and Dr. Josh McNamee. OpenHDR Viewer. Accessed July 7, 2024. https://viewer.openhdr.org/.
45. Marnerides, Demetris, Thomas Bashford-Rogers, Jonathan Hatchett, and Kurt Debattista. "Expandnet: A deep convolutional neural network for high dynamic range expansion from low dynamic range content." In Computer Graphics Forum, vol. 37, no. 2, pp. 37-49. 2018.
46. Zhu, Jun-Yan, Taesung Park, Phillip Isola, and Alexei A. Efros. "Unpaired image-to-image translation using cycle-consistent adversarial networks." In Proceedings of the IEEE international conference on computer vision, pp. 2223-2232. 2017.
47. Gardner, Marc-André, Kalyan Sunkavalli, Ersin Yumer, Xiaohui Shen, Emiliano Gambaretto, Christian Gagné, and Jean-François Lalonde. "Learning to predict indoor illumination from a single image." arXiv preprint arXiv:1704.00090 (2017).