

Few-Shot Domain Adaptation for Low Light RAW Image Enhancement

K Ram Prabhakar
ramprabhakar@iisc.ac.in
Vishal Vinod*
vishal114186@gmail.com
Nihar Ranjan Sahoo*
niharsahooigit@gmail.com
R. Venkatesh Babu
venky@iisc.ac.in

Video Analytics Lab,
Department of Computational and Data
Sciences,
Indian Institute of Science,
Bangalore, India

Abstract

Enhancing practical low light raw images is a difficult task due to severe noise and color distortions from short exposure time and limited illumination. Despite the success of existing Convolutional Neural Network (CNN) based methods, their performance is not adaptable to different camera domains. In addition, such methods also require large datasets with short-exposure and corresponding long-exposure ground truth raw images for each camera domain, which is tedious to compile. To address this issue, we present a novel few-shot domain adaptation method to utilize the existing source camera labeled data with few labeled samples from the target camera to improve the target domain's enhancement quality in extreme low-light imaging. Our experiments show that only ten or fewer labeled samples from the target camera domain are sufficient to achieve similar or better enhancement performance than training a model with a large labeled target camera dataset. To support research in this direction, we also present a new low-light raw image dataset captured with a Nikon camera, comprising short-exposure and their corresponding long-exposure ground truth images. The code is available at <https://val.cds.iisc.ac.in/HDR/BMVC21/index.html>.

1 Introduction

Capturing high-quality photos in low illumination is a fundamental yet challenging task. Increasing the ISO improves visibility; however, it also increases sensor noise. Longer exposure times improve the image but require a tripod to avoid camera motion and motion blur. Methods like enabling flash or image editing also present their own challenges. Low exposure image enhancement helps generate low-light scenes as if they were captured with a longer exposure time. It enables fast low illumination photography without a tripod. As shown in Fig. 1(a), the images captured in such settings possess a high degree of noise and color distortion. Existing single image denoising methods [8, 23] perform poorly and fail to correct color distortions in low light.

*Equal contribution



Figure 1: Qualitative comparison between different methods with an input to the model from the Canon dataset (left-most). The models are: LSID trained on (a) full 161 Sony source images, (b) only 6 Canon images, and (c) full Canon dataset. (d) 6 Canon images and 161 Sony images with proposed few-shot domain adaptation method (discussed in section 3).

An alternative is to merge a burst of short exposure images to reduce noise [12, 16, 20]. However, the burst images must be aligned for the camera and object motion, and aligning them in low-light conditions is a challenge. LSID, [6] a recent deep learning based method enhances raw low-light images by performing denoising, color improvement, and demosaicing, all with a single lightweight model. Despite the success of recent methods, there are two persisting challenges:

Domain Shift: CNNs are heavily data-sensitive; Since raw images captured with cameras from different manufacturers exhibit variations in color-space and noise characteristics, a model trained with one camera’s raw data performs sub-optimally on another camera’s raw data (see Table 1). Hence, there exists a domain shift across different camera raw domains. In this work, we consider each camera as a separate domain. As seen from the Fig. 1, cross-camera domain performance is poor as there are color distortions (green patches on the wall) and loss of finer details (missing cat’s whiskers) due to the shift across camera domains.

Tedious to collect labeled data: Collecting a large dataset of short-exposure and long-exposure raw image pairs for each camera is a difficult task. There must be no object motion, and to avoid camera misalignment, a tripod is necessary to capture long-exposure images (typically 10 seconds or more). Further, a smartphone or an IR remote is required to trigger the camera to avoid camera shake arising from physically pressing the camera click button. Thus, capturing a large-scale labeled dataset for different cameras is an arduous task.

To address the above mentioned challenges, we propose a paradigm shift for the low-light raw image enhancement task using few-shot learning and domain adaptation. We use a large collection of existing source camera labeled data to improve the performance and generate the output in the target domain by transferring the task onto a new target camera dataset with only a few labeled samples. In summary, our contributions are as follows:

- (i) To the best of our knowledge, we propose the first few-shot domain adaptation method for low-light raw image enhancement.
- (ii) We show that, with less than ten labeled samples from the target domain, our approach can outperform a model trained with a complete target domain dataset.
- (iii) We present experiments and ablations to illustrate the effectiveness of our method.
- (iv) We present a new Nikon camera dataset with short-exposure and long-exposure ground truth raw image pairs for the benefit of the research community.

2 Related Work

Low-light Image Enhancement: Existing low-light image enhancement methods require paired low/well-lit image scenes in RGB space and assume the images to be captured with minimal noise. While such methods capture global information suitably, their performance in extreme low-light conditions is sub-par. The histogram equalization method is useful for




Figure 2: First-order and second order statistics for an image captured with Nikon and Canon cameras. (a) Histogram of intensities, (b) Histogram of spatial derivatives of intensity, (c) Joint histogram of responses from convolution filters.

Table 1: Comparing LSID models trained with a complete source camera dataset and tested on respective target camera datasets.

Testing (→)	Sony		Nikon		Canon	
Training (↓)	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Sony [0]	28.50	0.774	25.90	0.693	27.41	0.845
Nikon	19.95	0.481	30.74	0.803	24.34	0.767
Canon [0]	18.51	0.542	23.27	0.847	32.32	0.899

Table 2: Details of the datasets used in our work.

Datasets	Exposure Ratios	Training Images	Testing Images
Sony [0]	90,15,300	161	36
Nikon	100,300	53	24
Canon [0]	50,150,300	44	21

increasing the dynamic range in a global context and is sub-optimal for extreme low-light enhancement [0]. Retinex [22] methods assume that the images have sufficient information to map the reflectance and enhance the low-light image. Similarly, [33] formulates the illumination estimation for enhancing underexposed images akin to expert retouched ground truth without image-to-image regression. Methods such as EnlightenGAN [24], a generative model with an attention U-Net [25], [37] for decomposition and enhancement and MIRNet [33] for attention aggregation have performed well in sRGB low-light enhancement.

Single Image Denoising: Non-deep methods for single image denoising either use engineered features [27, 29], or assume the noise model to be uniform and additive. Such assumptions in parametric methods are unsuitable for real-world low-light enhancement. Non-parametric methods depend on sparse image priors such as smoothness and self-similarity [8, 11, 18, 19] and are more expressive than parametric methods. BM3D [8, 23], a non-blind denoising method requiring noise and color model information, has outperformed deep methods in accuracy and noise robustness; however, it is prone to over smoothing in low light conditions. Several deep methods leveraging CNN based advancements have been proposed for denoising [13, 32, 36, 39, 40]. Autoencoders [17, 36] and MLPs [9] have shown sub-par performance on real-world raw sensor noise. Unprocess [8] learns the denoising pipeline and relevant photometric parameters by ‘unprocessing’ the image and is different from our goal to learn the camera’s color and noise model effectively.

Few-shot Domain Adaptation: Few-shot learning [10, 24, 30, 31] and Domain adaptation [0, 26] techniques are well explored in the context of many computer vision tasks. Several few-shot domain adaptation works [21, 28] use few labeled samples with many unlabeled samples in the target domain for image classification. Similarly, DA-FSL [43] is a few-shot domain adaptive prototypical learning method for recognition. The meta-learning paradigm [5] has also shown great promise in image denoising but depends on prior noise models to partially represent real noise. To the best of our knowledge, there has been no prior investigation of few-shot DA in inverse-imaging for raw camera domains.

As noted in [8], different camera sensors exhibit different noise models, and the process of capturing short-exposure and corresponding long-exposure raw images in low-light




Figure 3: (a) Example short-exposure and long-exposure image pairs from the Nikon dataset. The short exposure images are almost entirely dark whereas the long-exposure images have immense scene information. (b) Overview of our few-shot domain adaptation method.

conditions is expensive and time-consuming. While [34] has proposed efficient low-light enhancement, it is only for one type of smartphone camera. As a step toward tackling these challenges, we introduce the first of its kind few-shot domain adaptation and enhancement method for low-light conditions in the raw domain that is lightweight and highly effective.

3 Proposed method

With a noisy raw image captured with low-exposure time (i.e., shutter speed) as input, our CNN-based approach is trained to predict a clean long-exposure sRGB output of the same scene. The input is multiplied by an exposure factor calculated by the ratio of output and input exposure times. For example, to generate a 10-second long exposure output, the input 0.1-second low exposure image must be multiplied by 100. As a result of this operation, along with illumination, the noise is also amplified proportionally. Since we multiply the factor in the unprocessed raw domain and expect the output in the sRGB domain, the network must learn camera hardware-specific enhancement as well as its entire ISP pipeline (lens correction, demosaicing, white balancing, color manipulation, tone curve application, color space transform, and Gamma correction). Thus, a model trained on one specific camera data (source domain) does not translate similar performance to a different camera (target domain), hence the domain gap. In this paper, we propose to transfer the enhancement task from large labeled source data and generate output in the target domain using few labeled target data.

Problem formulation: We denote source domain (\mathbf{S}) with input short-exposure images as $\{S_n\}$ and corresponding long-exposure ground truth as $\hat{\mathbf{S}} = \{\hat{S}_n, \forall n = 1, \dots, N\}$. Similarly, the target domain (\mathbf{T}) consists of input images $\{T_m\}$ and corresponding ground truth, $\hat{\mathbf{T}} = \{\hat{T}_m, \forall m = 1, \dots, M\}$. Note that N is much greater than M , $N \gg M$. With both \mathbf{S} and \mathbf{T} as input, we train a CNN model (\mathcal{N}) to generate enhanced long-exposure output ($\tilde{\mathbf{S}}$ and $\tilde{\mathbf{T}}$). Our method is illustrated in Fig. 3(b) with the source and target training pipelines. It is an end-to-end trainable deep network that takes the raw sensor arrays as input and performs image enhancement utilizing the source data for few-shot domain adaptation to the target data.

Encoders: The significant domain gap between the source and target domains necessitates the extraction of separate and independent features from each domain before processing with a shared enhancement network (\mathcal{N}). Hence, we use a source encoder (\mathcal{E}_S) and a target encoder (\mathcal{E}_T). We first pack the input raw sensor arrays into a four-channel vector (for Bayer arrays from Sony, Nikon, and Canon cameras) and subtract the black level (reference volt-

age). Then, the packed array is multiplied by the exposure ratio factor and passed as input to the respective domain encoder. It should be noted that the exposure ratio factors need not be the same between the source and the target domain (See Table 2). For the encoder network, we use three convolutional layers with $\{16, 32, 64\}$ filters and 3×3 kernel size.

Enhancement Network: The source and target domain encoder features are passed separately to a shared common enhancement network, \mathbb{N} . By having a common enhancement network, the large pool of source data helps to improve the enhancement quality of \mathbb{N} , while the few target samples ensure that the output is in the target domain. We use U-Net architecture for the enhancement network. Further, the network has a pixel shuffle layer to convert 12-channel prediction to 16-bit three channel sRGB output. The objective of \mathbb{N} is to enhance, denoise, perform other ISP operations (AWB, color manipulation, etc.), and finally demosaicking to generate an sRGB output. \mathbb{N} generates enhanced output $\tilde{\mathbf{T}}$ for the target domain data as, $\tilde{\mathbf{T}} = \mathbb{N}(\mathcal{E}_T(\mathbf{T}))$. Similarly, $\tilde{\mathbf{S}}$ for the source domain as, $\tilde{\mathbf{S}} = \mathbb{N}(\mathcal{E}_S(\mathbf{S}))$.

Losses: For the target domain, we compute the ℓ_1 loss between the prediction ($\tilde{\mathbf{T}}$) and the ground truth ($\hat{\mathbf{T}}$) as, $\mathcal{L}_{target} = \ell_1(\tilde{\mathbf{T}}, \hat{\mathbf{T}})$. The source domain loss consists of two components: cosine similarity loss and SSIM loss. We compute cosine similarity between $\tilde{\mathbf{S}}$ and $\hat{\mathbf{S}}$ as, $\mathcal{L}_{CS}(\tilde{\mathbf{S}}, \hat{\mathbf{S}}) = 1 - \frac{\tilde{\mathbf{S}} \cdot \hat{\mathbf{S}}}{\|\tilde{\mathbf{S}}\| \times \|\hat{\mathbf{S}}\|}$. Cosine similarity loss is weak supervision for the source domain and is used instead of ℓ_1 loss since $N \gg M$, and using a strong supervision loss like ℓ_1 optimizes for pixel values to train \mathbb{N} , making the network predict the output in the source domain even for target domain input. Cosine similarity loss ensures that the prediction and the ground truth are in a similar direction. Hence with \mathcal{L}_{CS} , \mathbb{N} can still perform enhancement while predicting in target domain even for source domain input. Further, when trained with Sony as source and 4-shot Nikon as target (Table 7) with L_1 loss for the source, we obtain only 27.14dB PSNR for target domain validation, whereas using \mathcal{L}_{CS} loss for source achieves 30.30dB PSNR.

From experiments (in section 5), we find better enhancement (in terms of PSNR) using the structural similarity index measure (SSIM) [BS] to compute perceived degradation and preserve the spatial structure in the source output with respect to the ground truth. We do not use SSIM directly on the 16-bit data as that causes the source data to heavily influence the domain adaptation since the source dataset is much larger. Hence, we apply SSIM in JPEG compressed 8-bit domain, where the structural domain difference is less. Since type-casting the 16-bit data to 8-bit will still possess domain-specific details, we train a 16-to-8-bit U-net model (\mathcal{D} in Fig. 3) to convert the output from 16-bit to post-processed 8-bit representation.

The \mathcal{D} network is trained to perform the following non-linear operations: White balancing, Gamma correction, Quantization, and JPEG compression. Even after JPEG compression, the prediction may have traces of source domain specific color information. Further, the SSIM loss is a strong pixel-wise supervision, and in order to avoid the source domain from heavily influencing \mathbb{N} , we compute SSIM loss only in grayscale space, not in RGB color space. Also, it follows the intuition that the structure and edge information of a scene will remain the same across images captured with different cameras, while the color space representation may vary. We find that without SSIM loss for the source, we obtain 29.38dB PSNR on target domain validation, whereas using SSIM loss achieves 30.30dB PSNR (Table 7). For computing the SSIM loss, the ground truth ($\hat{\mathbf{S}}$) is also converted offline to post-processed 8-bit data ($\hat{\mathbf{S}}_{PP}$) using the rawpy post process function. Hence, the loss is obtained by computing SSIM loss between $\mathcal{D}(\tilde{\mathbf{S}})$ and $\hat{\mathbf{S}}_{PP}$, $\mathcal{L}_{SSIM} = 1 - SSIM(\mathcal{D}(\tilde{\mathbf{S}}), \hat{\mathbf{S}}_{PP})$. In Fig. 3, the top branch guided by the deep red arrows shows the entire source camera training

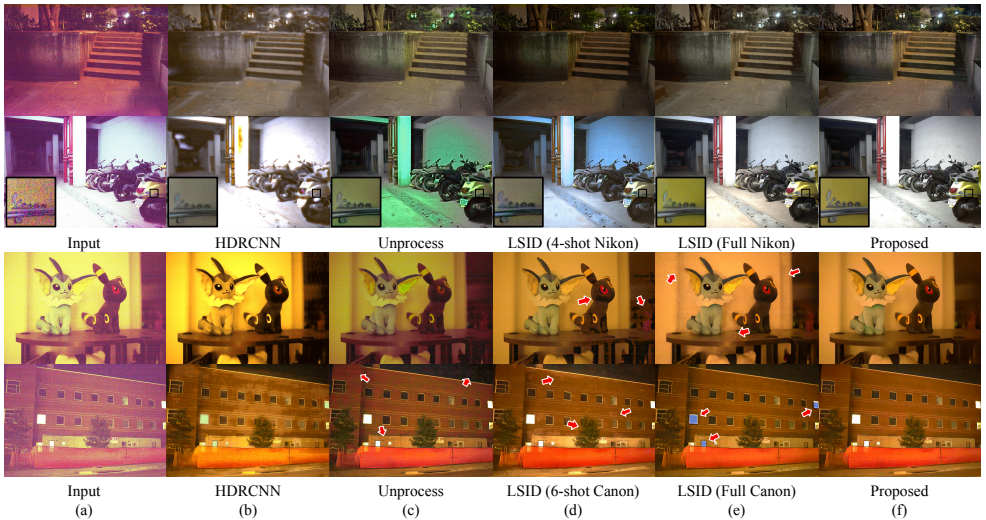


Figure 4: Qualitative comparison with methods tested on Nikon (top row) and Canon (bottom row) target images. (a) Input after multiplying by exposure factor, results from (b) HDRCNN and (c) Unprocess methods are after training on full Sony source and fine-tuning on k -shot target images, LSID with (d) k -shot target images and (e) full target training dataset ($k=53$ for Nikon and $k=44$ for Canon), (f) Proposed few-shot domain adaptation method with 161 Sony source images and 4-shot Nikon (top row) and 6-Canon (bottom row) target images.

pipeline. It should be noted that \mathcal{D} is used only to compute the loss but not in inference. Finally, we use the sum of cosine similarity loss (\mathcal{L}_{CS}) as well as the SSIM loss calculated in the 8-bit domain as the total loss for the source camera pipeline: $\mathcal{L}_{source} = \mathcal{L}_{CS} + \mathcal{L}_{SSIM}$. The total loss is the sum of target and source domain losses: $\mathcal{L}_{total} = \mathcal{L}_{target} + \mathcal{L}_{source}$.

4 Experiments

Datasets: We use the Sony camera dataset [9] for our source training pipeline. We expect the diverse, high-quality low-light scenes from this dataset to aid few-shot domain adaptation performance in terms of the color spaces and noise model learned by our method.

For few-shot domain adaptation, we work with very few (< 10) target camera images in every training experiment. We use the open-source Canon camera low-light raw image dataset [10] and a new Nikon camera dataset that we have compiled and make available with this work for our target camera training pipeline. We do not investigate burst denoising or the ‘lucky imaging’ phenomenon. Hence, we only take the first short-exposure raw image for each scene from the Sony dataset and use the 161 images for our source camera training pipeline. Note that the Canon dataset has eight different ratios with close ranges such that they can be put into three buckets of ratios: 50, 150, and 300 (Table 2).

Nikon camera dataset: We have compiled a dataset of raw low-light images captured with a Nikon D5600 camera to train the proposed few-shot domain adaptation architecture. The Nikon dataset consists of short-exposure images captured at $1/3$ or $1/10$ seconds and corresponding ground-truth long-exposure images captured at 10 or 30 seconds in the NEF format. For uniformity, there are two short-exposure images for every long-exposure image

Table 3: Quantitative comparison of Sony as source with Nikon or Canon as target. The improvement of our method over only k shot trained model is in brackets. The LSID model trained with full Nikon dataset ($k=53$) achieves 30.74dB PSNR and 0.803 SSIM and when trained with full Canon dataset ($k=44$) attains 32.32dB PSNR and 0.899 SSIM (see Table 1).

Nikon as target				SSIM		
k (\rightarrow)	1	2	4	1	2	4
LSID (only k target)	23.20 \pm 3.06	27.27 \pm 0.384	28.05 \pm 1.53	0.679 \pm 0.172	0.819 \pm 0.031	0.864 \pm 0.011
Proposed (k target + source)	25.27 \pm 0.58 (+2.07)	28.06 \pm 0.671 (+0.79)	30.30 \pm 0.52 (+2.25)	0.860 \pm 0.010 (+0.181)	0.909 \pm 0.003 (+0.090)	0.913 \pm 0.006 (+0.049)
Canon as target				SSIM		
k (\rightarrow)	1	3	6	1	3	6
LSID (only k target)	21.54 \pm 2.89	26.9 \pm 2.37	29.36 \pm 0.763	0.588 \pm 0.182	0.785 \pm 0.005	0.829 \pm 0.007
Proposed (k target + source)	24.29 \pm 3.16 (+2.75)	28.78 \pm 3.54 (+1.8)	33.22 \pm 0.45 (+3.86)	0.623 \pm 0.007 (+0.035)	0.841 \pm 0.033 (+0.056)	0.896 \pm 0.015 (+0.067)

such that the exposure ratio (ratio of exposure time between the ground-truth long-exposure image and the input short-exposure image) is 100 and 300, respectively. Similar to [6], we mount the camera on sturdy tripods and use appropriate camera settings to capture the static scenes using a smartphone app. The images captured include 129 short-exposure and 65 long-exposure ground-truth images of indoor and outdoor low-light scenes (sub lux).

Training Setup: We train the source and target pipelines simultaneously in an end-to-end manner. As discussed in section 3, we use the respective short-exposure raw images as the input to each of the encoders. We first randomly crop a 512×512 image patch and augment it with random-flip and random-rotate. We then subtract the black level and multiply the input raw image with the exposure ratio. We use an initial learning rate of 10^{-4} up to 2000 epochs and then reduce it by a factor of 10 for every 1000 epochs thereafter. We use the Adam [14] optimizer for the 8-bit SSIM loss and Cosine Similarity loss (\mathcal{L}_{source}) for the source pipeline, and the \mathcal{L}_{target} loss for the target pipeline.

We train the model for 4000 epochs (same as [6]) but observe the loss saturating at lower epochs prompting us to employ early stopping. Since the large source domain has 161 images, every epoch has 161 train steps. As we jointly train the source and target pipelines, for every epoch, we use 161 randomly cropped source patches obtained from the 161 source domain raw images and 161 randomly cropped target patches from only k -images in the target domain. We find that training our proposed method for up to 2500 epochs is sufficient to obtain the best results and reproduce the results in this paper for few-shot domain adaptation.

The source SSIM loss is calculated in the 8-bit space after passing the output from the shared \mathbb{N} through the 16-to-8-bit converter. The cosine similarity loss for the source domain and L_1 loss for the target domain are computed in the 16-bit sRGB space. The exposure ratio is computed and provided to the network. At inference time, we use the full-scale raw target image as input to the target camera pipeline and obtain the enhanced target sRGB image.

5 Results

Table 2 lists the total number of labeled raw image pairs in the train set to be 161 images for the Sony dataset [6], 53 images for our Nikon dataset, and 44 images for the Canon dataset [10]. We train our model for three different numbers of labeled target data: For Nikon target, we use $k=1,2$, and 4. For Canon target, we use $k=1,3$, and 6. We choose the k based on the different exposure ratios available for the dataset. As reported in Table 2, the Nikon




Figure 5: Qualitative comparison of results for $k=1,2$ and 4 for Sony as source and Nikon as target, and $k=1,3$ and 6 for Sony as source and Canon as target for choosing the value of k .

Table 4: Comparison with baselines trained with Sony raw data as source and fine-tuned with (1/2/4-shot) Nikon camera raw images and (1/3/6-shot) Canon camera raw images. We compare with HDRCNN [9], Unprocess [9] and LSID [9].

Method	Sony source w/ Nikon target						Sony source w/ Canon target					
	1		2		4		1		3		6	
$k \rightarrow$	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
HDRCNN	11.47	0.648	12.36	0.578	12.87	0.627	15.07	0.593	15.89	0.627	16.14	0.633
Unprocess	18.90	0.710	21.91	0.690	25.63	0.761	16.97	0.609	22.15	0.671	27.20	0.733
LSID	22.38	0.746	25.75	0.874	27.93	0.899	25.88	0.792	28.38	0.831	28.85	0.819
Proposed	25.27	0.860	28.06	0.909	30.30	0.913	24.29	0.623	28.78	0.841	33.22	0.896

camera dataset has two exposure ratios; hence we use two images per ratio leading to a total of four images. Similarly, the Canon camera dataset has three ratios; hence we use two images per ratio leading to a total of six images. We observe that using two images per exposure ratio in the target domain is sufficient to outperform all baselines (Refer Table 3). A qualitative guideline for choosing the value of k is in Fig.(5). For each k , we run three separate experiments, each with a different set of k labeled target images. We report the average and 95% variance margin computed across three different sets for the Nikon dataset and for the Canon dataset (Table 3). We compare with an LSID model trained with only k -target data and an LSID model trained with the full target camera data.

Quantitative Evaluation: For Nikon dataset as target, our 4-shot approach achieves 30.30dB PSNR, which is on par with the full target dataset ($k=53$) trained LSID model. Our method outperforms only k -shot trained model by 2.25dB PSNR (Table 3). Similarly, for Canon dataset as target, our 6-shot approach outperforms the full dataset ($k=44$) trained LSID model by 0.9dB PSNR, and k -shot trained LSID model by 3.86dB PSNR (Table 3). Since ours is the first few-shot domain adaptation method, we quantitatively compare our method with baselines by training them on the full Sony source dataset and then fine-tuning them in a few-shot manner on the Nikon or Canon target datasets (Table 4). We also compare with recent low-light enhancement methods in Table 5 and show that our method outperforms all k -shot fine-tuned baselines in PSNR and SSIM for the Nikon and Canon target datasets. We show quantitative results for our method trained with Sony as source and four OnePlus camera images or four Google Pixel camera images as target in Table 6.

As discussed in section 1, capturing a low-light raw image dataset is difficult, and different cameras have different color-space and noise distributions, hence there is a need for a domain adaptation method that can transfer the task from source to a target domain in a few-shot setting. Despite the high complexity of the task, our method outperforms all base-

Table 5: Comparison with recent low-light image enhancement methods. Ours is the first few-shot domain adaptation method, hence, we first train the baselines on the Sony camera dataset and then fine-tune with (4-shot) Nikon camera or (6-shot) Canon camera datasets.

Method	Sony Source		Sony w/ Nikon		Sony w/ Canon	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
DeepUPE [13]	14.58	0.256	13.42	0.266	13.81	0.285
MIRNet [65]	15.24	0.414	14.18	0.458	13.24	0.397
KnD [10]	17.15	0.313	15.04	0.226	17.24	0.432
HDRCNN [9]	17.39	0.491	12.87	0.627	16.14	0.633
KnD++ [10]	23.03	0.579	19.38	0.471	21.25	0.434
Unprocess [9]	27.83	0.700	25.63	0.761	27.20	0.733
LSID [9]	28.50	0.774	27.93	0.899	29.36	0.829
Proposed	-	-	30.30	0.913	33.22	0.896

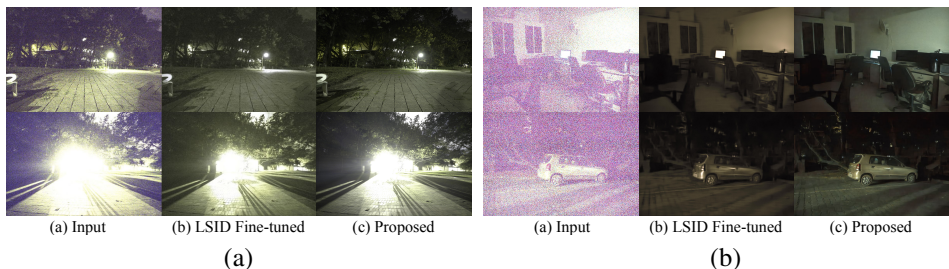


Figure 6: Results for Sony as source with (a) OnePlus 5 as target (b) Google Pixel as target.

lines with a lightweight model because the abundant source data helps to learn the low-light enhancement task in the source+ k -shot setting successfully as compared to using only target domain data. While the k -shot samples help to predict the output in the target domain, the task is transferred successfully from the large source domain to the target camera domain.

Qualitative Evaluation: In Fig. 4, we qualitatively compare the results from our method and baselines when trained with Sony as source and Nikon (top row) or Canon (bottom row) as target. As highlighted by the zoomed-in regions and red arrows, the baseline results have several artifacts in terms of noise and color. Although the LSID model trained with full target data performs better than the k -shot model (for Nikon or Canon as target), it is still sub-par compared to our method’s results. We also show qualitative results for our method trained with Sony as source and four OnePlus camera images (Fig. 6(a)) or four Google Pixel camera images (Fig. 6(b)) as target to demonstrate the effectiveness of our method on low-cost smartphone camera data, which typically have higher noise severity in low-light.

Ablation Study: We discuss relevant ablations for our proposed method trained with Sony as source and 4-shot Nikon as target (Refer Table 7). The details are as follows:

- *Separate Encoder and Decoder:* Transfer learning methods performed sub-optimally primarily due to the color bias. Hence, we trained separate encoders, a shared U-net, and separate camera-specific decoders to learn different camera color spaces individually. This approach obtained a PSNR of 28.62dB and SSIM of 0.867, and the results had visible color gaps between the model’s output and ground truth.
- *Combined Encoders:* To verify camera-specific denoising by the encoders, we used a shared encoder followed by a shared \mathbb{N} network. But, we noticed the colors of the target domain’s output to be dominated by the large source domain data. We obtain a PSNR of 29.20dB and SSIM of 0.890, which is better than fine-tuning LSID, but the colors of the target domain’s output validation images were dominated by the source training images.
- *Proposed w/ Source ℓ_1 :* Training with a strong supervision loss for the source domain

Table 6: Quantitative evaluation of Sony as source with OnePlus or Pixel as target. The LSID model trained with full OnePlus dataset ($k=12$) achieves 28.88dB PSNR and 0.708 SSIM and when trained with full Pixel dataset ($k=10$) attains 27.95dB PSNR and 0.759 SSIM. (Note: test set for OnePlus=38 images and Pixel=18 images).

Method	Sony source w/ OnePlus target						Sony source w/ Pixel target					
	1		2		4		1		2		4	
k (\rightarrow)	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LSID	22.54	0.617	25.87	0.651	26.69	0.670	13.86	0.296	21.91	0.695	27.62	0.760
Proposed	25.10	0.633	26.47	0.650	27.77	0.712	18.60	0.635	24.98	0.753	29.60	0.782

Table 7: Ablation study for training with Sony source and 4-shot Nikon as target dataset. The table reports mean PSNR/SSIM over three different runs. See section 5 for details.

No.	Details	PSNR	SSIM
1	Separate encoder and decoder	28.62	0.867
2	Combined encoder	29.20	0.890
3	Proposed w/ Source ℓ_1 loss	27.14	0.807
4	Proposed w/o Source SSIM loss	29.38	0.902
5	Proposed	30.30	0.913

strongly influenced the colors of the target domain’s output images.

- *Proposed w/o \mathcal{L}_{SSIM}* : From previous ablations, we observe the source domain influencing the target output image color. Hence, we separate only the encoder while merging the U-net and the decoder to get (\mathbb{N}). We experimented with several loss functions, including combinations of ℓ_2 loss, grayscale SSIM, gradient loss, ℓ_1 loss, and cosine similarity loss. We found that using cosine similarity loss and SSIM loss for the source domain, and ℓ_1 loss [24] for the target led to better preservation of color and structure information.

- *The \mathcal{D} network for SSIM loss*: As discussed in section 3, type-casting the source output from 16-bit to 8-bit space will still possess domain specific details. The SSIM loss is used to compare the brightness and structural details but not the color quality (cosine similarity is for color). Thus, following the camera ISP processing steps, we train a 16-to-8-bit conversion U-net model (\mathcal{D} in Fig. 2(b)) to convert the 16-bit data to 8-bit data. The \mathcal{D} network is trained to perform the following non-linear operations: White balancing, Gamma correction, Quantization and JPEG compression. From experiments, we observe that a U-net is necessary to learn all the above mentioned non-linear operations.

6 Conclusion

In this work, we propose a novel domain adaptation method for low-light raw image enhancement using only a few labeled samples from the target domain and many source domain samples. We first extract camera-specific information through separate encoders, and then use a shared enhancement network (\mathbb{N}) to extract domain invariant features and transfer the task successfully from the source domain to the target domain. We propose to compute post-processed image loss for the source domain to have less influence on predictions in the target domain. We also present a new labeled raw image dataset captured with a Nikon camera. Our results show that using only a few labeled samples from the target domain is sufficient to obtain similar or better results than training with large target domain data. We hope that our method and dataset inspire new investigations along these research directions.

Acknowledgment: This work was supported by a project grant from MeitY (No.4(16)/2019-ITEA), Govt. of India.

References

- [1] Canon 6D datasets for learning to see in the dark. <https://github.com/jconenna/Canon-6D-Datasets-For-Learning-to-See-in-the-Dark>. Online; Accessed: 2020-09-30.
- [2] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [3] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T. Barron. Unprocessing images for learned raw denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [4] Harold C Burger, Christian J Schuler, and Stefan Harmeling. Image denoising: Can plain neural networks compete with BM3D? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2012.
- [5] Leslie Casas, Attila Klimmek, Gustavo Carneiro, Nassir Navab, and Vasileios Belagiannis. Few-shot meta-denoising. *arXiv preprint arXiv:1908.00111*, 2019.
- [6] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [7] Jingwen Chen, Jiawei Chen, Hongyang Chao, and Ming Yang. Image blind denoising with generative adversarial network based noise modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [8] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8):2080–2095, 2007.
- [9] Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafal Mantiuk, and Jonas Unger. HDR image reconstruction from a single exposure using deep CNNs. *ACM Transactions on Graphics (TOG)*, 36(6):1–15, 2017.
- [10] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017.
- [11] Shuhang Gu, Lei Zhang, Wangmeng Zuo, and Xiangchu Feng. Weighted nuclear norm minimization with application to image denoising. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014.
- [12] Samuel W Hasinoff, Dillon Sharlet, Ryan Geiss, Andrew Adams, Jonathan T Barron, Florian Kainz, Jiawen Chen, and Marc Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics (TOG)*, 35(6):1–12, 2016.
- [13] Viren Jain and Sebastian Seung. Natural image denoising with convolutional networks. *Advances in neural information processing systems*, 21:769–776, 2008.

- [14] Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou, and Zhangyang Wang. Enlightengan: Deep light enhancement without paired supervision. *IEEE Transactions on Image Processing*, 30:2340–2349, 2021.
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] Ziwei Liu, Lu Yuan, Xiaoou Tang, Matt Uyttendaele, and Jian Sun. Fast burst images denoising. *ACM Transactions on Graphics (TOG)*, 33(6):1–9, 2014.
- [17] Kin Gwn Lore, Adedotun Akintayo, and Soumik Sarkar. Llnet: A deep autoencoder approach to natural low-light image enhancement. *Pattern Recognition*, 61:650–662, 2017.
- [18] Artur Łoza, David R Bull, Paul R Hill, and Alin M Achim. Automatic contrast enhancement of low-light images based on local statistics of wavelet coefficients. *Digital Signal Processing*, 23(6):1856–1866, 2013.
- [19] Julien Mairal, Francis Bach, Jean Ponce, Guillermo Sapiro, and Andrew Zisserman. Non-local sparse models for image restoration. In *IEEE 12th International Conference on Computer Vision*, September 2009.
- [20] Ben Mildenhall, Jonathan T Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. Burst denoising with kernel prediction networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [21] Saeid Motiian, Quinn Jones, Seyed Mehdi Iranmanesh, and Gianfranco Doretto. Few-shot adversarial domain adaptation. *arXiv preprint arXiv:1711.02536*, 2017.
- [22] Seonhee Park, Soohwan Yu, Byeongho Moon, Seungyong Ko, and Joonki Paik. Low-light image enhancement using variational optimization-based retinex model. *IEEE Transactions on Consumer Electronics*, 63(2):178–184, 2017.
- [23] Tobias Plotz and Stefan Roth. Benchmarking denoising algorithms with real photographs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [24] K Ram Prabhakar, Gowtham Senthil, Susmit Agrawal, R Venkatesh Babu, and Rama Krishna Sai S Gorthi. Labeled from unlabeled: Exploiting unlabeled data for few-shot deep HDR dehazing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [25] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [26] Artem Rozantsev, Mathieu Salzmann, and Pascal Fua. Beyond sharing weights for deep domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(4):801–814, 2018.
- [27] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992.

- [28] Doyen Sahoo, Hung Le, Chenghao Liu, and Steven CH Hoi. Meta-learning with domain adaptation for few-shot learning under domain shift. 2018.
- [29] Eero P Simoncelli and Edward H Adelson. Noise removal via bayesian wavelet coring. In *Proceedings of 3rd IEEE International Conference on Image Processing*, volume 1, pages 379–382. IEEE, 1996.
- [30] Qianru Sun, Yaoyao Liu, Tat-Seng Chua, and Bernt Schiele. Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [31] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [32] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [33] Ruixing Wang, Qing Zhang, Chi-Wing Fu, Xiaoyong Shen, Wei-Shi Zheng, and Jiaya Jia. Underexposed photo enhancement using deep illumination estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [34] Yuzhi Wang, Haibin Huang, Qin Xu, Jiaming Liu, Yiqun Liu, and Jue Wang. Practical deep raw image denoising on mobile devices. In *European Conference on Computer Vision (ECCV)*. Springer, 2020.
- [35] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [36] Junyuan Xie, Linli Xu, and Enhong Chen. Image denoising and inpainting with deep neural networks. In *Advances in Neural Information Processing Systems*, pages 341–349, 2012.
- [37] Ke Xu, Xin Yang, Baocai Yin, and Rynson WH Lau. Learning to restore low-light images via decomposition-and-enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [38] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for real image restoration and enhancement. In *Proceedings of the European Conference on Computer Vision (ECCV)*, August 2020.
- [39] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017.
- [40] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for CNN-based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018.

- [41] Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. Kindling the darkness: A practical low-light image enhancer. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 1632–1640, 2019.
- [42] Yonghua Zhang, Xiaojie Guo, Jiayi Ma, Wei Liu, and Jiawan Zhang. Beyond brightening low-light images. *International Journal of Computer Vision*, 129(4):1013–1037, 2021.
- [43] An Zhao, Mingyu Ding, Zhiwu Lu, Tao Xiang, Yulei Niu, Jiechao Guan, and Ji-Rong Wen. Domain-adaptive few-shot learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2021.
- [44] Hang Zhao, Orazio Gallo, Iuri Frosio, and Jan Kautz. Loss functions for image restoration with neural networks. *IEEE Transactions on Computational Imaging*, 3(1):600–612, 2017.