

RAW-Adapter: Adapting Pre-trained Visual Model to Camera RAW Images

Ziteng Cui¹ and Tatsuya Harada^{1,2}

¹ The University of Tokyo

² RIKEN AIP

{cui, harada}@mi.t.u-tokyo.ac.jp

Abstract. sRGB images are now the predominant choice for pre-training visual models in computer vision research, owing to their ease of acquisition and efficient storage. Meanwhile, the advantage of RAW images lies in their rich physical information under variable real-world challenging lighting conditions. For computer vision tasks directly based on camera RAW data, most existing studies adopt methods of integrating image signal processor (ISP) with backend networks, yet often overlook the interaction capabilities between the ISP stages and subsequent networks. Drawing inspiration from ongoing adapter research in NLP and CV areas, we introduce **RAW-Adapter**, a novel approach aimed at adapting sRGB pre-trained models to camera RAW data. RAW-Adapter comprises input-level adapters that employ learnable ISP stages to adjust RAW inputs, as well as model-level adapters to build connections between ISP stages and subsequent high-level networks. Additionally, RAW-Adapter is a general framework that could be used in various computer vision frameworks. Abundant experiments under different lighting conditions have shown our algorithm’s state-of-the-art (SOTA) performance, demonstrating its effectiveness and efficiency across a range of real-world and synthetic datasets. Code is available at [this url](#).

1 Introduction

In recent years, there has been a growing interest in revisiting vision tasks using unprocessed camera RAW images. How to leverage information-rich RAW image in computer vision tasks has become a current research hotspot in various sub-areas (*i.e.* denoising [2, 76], view synthesis [46], object detection [27, 48]). Compared to commonly used sRGB image, RAW image directly acquired by the camera sensor, encompasses abundant information unaffected or compressed by image signal processor (ISP), also offers physically meaningful information like noise distributions [47, 68], owing to its linear correlation between image intensity and the radiant energy received by a camera. The acquisition of RAW data, facilitating enhanced detail capture and a higher dynamic range, imparts a unique advantage in addressing visual tasks under variable lighting conditions in the real world. For instance, sunlight irradiance can reach as high as $1.3 \times 10^3 W/m^2$, while bright planet irradiance can be as low as $2.0 \times 10^{-6} W/m^2$ [32].




Fig. 1: (a). An overview of basic image signal processor (ISP) pipeline. (b). ISP and current visual model have different objectives. (c) Previous methods optimize ISP with down-stream visual model. (d) Our proposed RAW-Adapter.

Meanwhile, sRGB has emerged as the primary choice for pre-training visual models in today’s computer vision field, due to its scalability and ease of storage. Typically, sRGB images are derived from camera RAW data through the ISP pipeline. As shown in Fig. 1(a), the entire ISP pipeline consists of multiple modules to convert RAW images into vision-oriented sRGB images, each of these modules serves its own distinct purpose, with the majority being ill-posed and heavily dependent on prior information [5, 16, 26, 35, 42, 56].

When adopt RAW data for computer vision tasks, the purpose of a manually designed ISP is to produce images that offer a superior visual experience [18, 69], rather than optimizing for downstream visual tasks such as object detection or semantic segmentation (see Fig. 1(b)). Additionally, most companies’ ISPs are black boxes, making it difficult to obtain information about the specific steps inside. Consequently, utilizing human vision-oriented ISP in certain conditions is sometimes even less satisfactory than directly using RAW [22, 27, 40, 81].

To better take advantage of camera RAW data for various computer vision tasks. Researchers began to optimize the image signal processor (ISP) jointly with downstream networks. Since most ISP stages are non-differentiable and cannot be jointly backpropagation, there are two main lines of approaches to connect ISP stages with downstream networks (see Fig. 1(c)): ①. First-kind approaches maintain the modular design of the traditional ISP, involving the design of differentiable ISP modules through optimization algorithms [48, 69, 75], such as Hardware-in-the-loop [48] adopt covariance matrix adaptation evolution strategy (CMA-ES) [23]. ②. Second-kind approaches replaced the ISP part entirely with a neural network [18, 55, 72, 73], such as Dirty-Pixel [18] adopt a stack of residual UNets [58] as pre-encoder, however, the additional neural network introduces a significant computational burden, especially when dealing with high-resolution inputs. Beyond these, the above-mentioned methods still treat ISP and the backend neural network as two independent modules, lacking interac-




Fig. 2: Performance of RAW-based visual tasks with and without sRGB pre-trained weights. We analyze two methods: Dirty-Pixel [18] and RAW-Adapter. Blue line represents trained with MS COCO [44] pre-train weights, the purple line indicates ImageNet [17] pre-train weights, and the yellow line signifies training from scratch.

tion capabilities between two separate models. Especially, current visual models predominantly pre-trained on large amounts of sRGB images, there exists a significant gap between pre-trained models and camera RAW images. As it shown in Fig. 2, training from scratch with RAW data can significantly impair the performance of RAW-based vision tasks. How to leveraging sRGB pre-trained weights proves to be highly beneficial for RAW image visual tasks. Hence, here we ask: Is there a more effective way to adapt the **information-rich** RAW data into the **knowledge-rich** sRGB pre-trained models?

Our solution, RAW-Adapter, differs from previous approaches that employed complex ISP or deeper neural networks at the input level. Instead, we prioritize simplifying the input stages and enhancing connectivity between ISP and subsequent networks at the model level. Inspired by recent advancements in prompt learning and adapter tuning [8, 28, 33, 54], we developed two novel adapter approaches aimed at enhancing the integration between RAW input and sRGB pre-trained models. Leveraging priors from ISP stages, our method includes input-level adapters and model-level adapters, as illustrated in Fig. 1(d). Input-level adapters are designed to adapt the input from RAW data to the backend network input, we follow traditional ISP design and predict ISP stages’ key parameters to accommodate RAW image into the specific down-stream vision tasks, in this part, we employ strategies involving Query Adaptive Learning (QAL) and Implicit Neural Representation (INR) to ensure the differentiability of the core ISP processes while also upholding a lightweight design. Model-level adapters leverage prior knowledge from the input-level adapters by extracting intermediate stages as features, these features are then embedded into the downstream network, thus the network incorporates prior knowledge from the ISP stages, then jointly contributes to final machine-vision perception tasks. Our contributions could be summarized as follows:

- We introduce RAW-Adapter, a novel framework aimed at improving the integration of sRGB pre-trained models to camera RAW images. Through the implementation of two types of adapters, we strive to narrow the dispar-

ity between RAW images and sRGB models, tackling discrepancies at both input and model levels.

- By analyzing camera ISP steps, we’ve crafted input-level adapters that integrate query adaptive learning (QAL) and implicit neural representations (INR) to optimize ISP key parameters. Additionally, harnessing prior input stage information informs model-level adapters, enriching model understanding and improving downstream task performance.
- We conducted detection and segmentation experiments involving different lighting scenarios, including normal-light, dark, and over-exposure scenes. Through comparisons with current mainstream ISPs and joint-training methods, sufficient experiments demonstrate that our algorithm achieved state-of-the-art (SOTA) performance.

2 Related Works

2.1 Image Signal Processor

Typically, a camera’s Image Signal Processor (ISP) pipeline is required to reconstruct a high-quality sRGB image from camera RAW data, traditional ISP pipeline is formulated as a series of manually crafted modules executed sequentially [16, 35, 51, 56], including some representative steps such as demosaicing, white balance, noise removal, tone mapping, color space transform and so on. There also exist alternative designs of ISP process such as Heide *et al.* [26] designed FlexISP which combine numerous ISP blocks to a joint optimization block, and Hasinoff *et al.* [24] modified some of the traditional ISP steps for burst photography under extreme low-light conditions. When it comes to deep learning era, Chen *et al.* [6] proposed SID to use a UNet [58] replace the traditional ISP steps, which translate low-light RAW data to normal-light sRGB images, and Hu *et al.* [29] proposed a white-box solution with eight differentiable filters. After that, many efforts began to substitute the traditional ISP process with a neural network, such as [6, 10, 30, 34, 37, 42, 60, 62, 79], however deep network models are constrained by the training dataset, leading to shortcomings in generalization performance. Meanwhile another line of work focus on translate sRGB images back to RAW data [2, 39, 49, 50, 67, 71, 76]. Apart from the conversion between RAW and RGB, there is ongoing research specifically dedicated to harnessing RAW images for downstream computer vision tasks.

2.2 Computer Vision based on RAW data

In order to effectively leverage information in RAW images for downstream visual tasks and save the time required for ISP, some early methods proposed to directly perform computer vision tasks on RAW images [3, 81], which lacks consideration for the camera noise introduced during the conversion from photons to RAW images, especially in low-light conditions [41, 47, 68]. Meanwhile training from scratch on RAW data would abandon the current large-scale pre-trained

visual model on sRGB data (see Fig. 2), especially the quantity of RAW image datasets (*i.e.* 4259 images in PASCAL RAW [52] dataset and 2230 images in LOD [27] dataset) is incomparable to current RGB datasets (*i.e.* over 1M images in ImageNet [17] dataset and SAM [38] is trained with 11M images).

Most subsequent research focuses on finding ways to integrate ISP and backend computer vision models [18, 22, 48, 55, 69, 72, 73, 75]. For example Wu *et al.* [69] first propose VisionISP and emphasized the difference between human vision and machine vision, they introduced several trainable modules in the ISP to enhance backend object detection performance. Since then several methods attempt to replace the existing non-differentiable ISP with a differentiable ISP network [48, 75] or directly using encoder networks to replace the ISP process [18, 55, 73], however training two consecutive networks simultaneously can result in a significant consumption of computational resources, previous work also rarely focused on how to fine-tune the current mainstream sRGB visual models more efficiently on RAW data. Additionally, research like Rawgment [74], focusing on achieving realistic data augmentation directly on RAW images, some studies address the decrease in computer vision performance caused by specific steps within the in-camera process or ISP, such as white balance error [1], auto-exposure error [53], camera motion blur [59] and undesirable camera noise [15, 21].

2.3 Adapters in Computer Vision

Adapters have become prevalent in NLP area, which introducing new modules in large language model (LLM) such as [28, 63], these modules enable task-specific fine-tuning, allowing pre-trained models to swiftly adapt to downstream NLP tasks. In the computer vision area, adapters have been adopted in various areas such as incremental learning [31] and domain adaptation [57]. Recently, a series of adapters have been used to investigate how to better utilize pre-trained visual models [8, 33, 36] or pre-trained vision-language models [65, 77]. These methods focus on utilizing prior knowledge to make pre-trained models quickly adapt to downstream tasks. Here, we introduce the RAW-Adapter to further improve the alignment between RAW inputs and sRGB pre-trained visual models.

3 RAW-Adapter

In this section, we introduce RAW-Adapter. Due to page limitation, an overview of the conventional image signal processor (ISP) in supplementary’s Sec. A. For specific adapter designs, we selectively omit certain steps and focus on some key steps within the camera ISP. In Sec. 3.1 we introduce RAW-Adapter’s input-level adapters, followed by an explanation of its model-level adapters in Sec. 3.2.

3.1 Input-level Adapters

Fig. 3(a) provides an overview of RAW-Adapter. The solid lines in Fig. 3(a) left represent input-level adapters. Input-level adapters are designed to convert




Fig. 3: (a). Structure of RAW-Adapter. Solid line in left denotes input-level adapter’s workflow and dotted line denotes model-level adapter’s workflow, stage 1~4 means different stage of visual model backbone. (b). Detailed structure of kernel & matrix predictors P_K, P_M . (c). Detailed structure of model-level adapter M ’s merge block.

the RAW image I_1 into the machine-vision oriented image I_5 . This process involves digital gain & denoise, demosaicing, white balance adjustment, camera color matrix, and color manipulation.

We maintain the ISP design while simultaneously using query adaptive learning (QAL) to estimate key parameters in ISP stages. The QAL strategy is motivated by previous transformer models [4, 9, 13] and detailed structure is shown in Fig. 3(b), input image $I_{i \in (1,2)}$ would pass by 2 down-scale convolution blocks to generate feature, then feature pass by 2 linear layers to generate attention block’s key k and value v , while query q is a set of learnable dynamic parameters, the ISP parameters would be predicted by self-attention calculation [66]:

$$parameters = FFN(\text{softmax}(\frac{q \cdot k^T}{\sqrt{d_k}}) \cdot v), \quad (1)$$

where FFN denotes the feed-forward network, includes 2 linear layers and 1 activation layer, the predicted parameters would keep the same length as query q . We defined 2 QAL blocks P_K and P_M to predict different part parameters.

The input RAW image I_1 would first go through the pre-process operations and a demosaicing stage [16, 35], followed by subsequent ISP processes:

Gain & Denoise: Denoising algorithms always take various factors into account, such as input photon number, ISO gain level, exposure settings. Here we

first use QAL block $\mathbb{P}_{\mathbb{K}}$ to predict a gain ratio g [2,6] to adapt \mathbf{I}_1 in different lighting scenarios, followed by an adaptive anisotropic Gaussian kernel to suppress noise under various noise conditions, $\mathbb{P}_{\mathbb{K}}$ will predict the appropriate Gaussian kernel k for denoising to improve downstream visual tasks' performance, the predicted key parameters are the Gaussian kernel's major axis r_1 , minor axis r_2 (see Fig. 4 left), here we set the kernel angle θ to 0 for simplification. Gaussian kernel k at pixel (x, y) would be:

$$k(x, y) = \exp(-(b_0x^2 + 2b_1xy + b_2y^2)), \quad (2)$$

where:

$$b_0 = \frac{\cos(\theta)^2}{2r_1^2} + \frac{\sin(\theta)^2}{2r_2^2}, b_1 = \frac{\sin(2\theta)}{4r_1^2} \left(\left(\frac{r_1}{r_2} \right)^2 - 1 \right), b_2 = \frac{\sin(\theta)^2}{2r_1^2} + \frac{\cos(\theta)^2}{2r_2^2}. \quad (3)$$

After gain ratio g and kernel k process on image \mathbf{I}_1 , $\mathbb{P}_{\mathbb{K}}$ also predict a filter parameter σ (initial at 0) to keep the sharpness and recover details of generated image \mathbf{I}_2 . Eq. 4 shows the translation from \mathbf{I}_1 to \mathbf{I}_2 , where \otimes denotes kernel convolution and filter parameter σ is limited in a range of (0, 1) by a Sigmoid activation. For more details please refer to our supplementary part Sec. C.

$$\begin{aligned} \mathbb{P}_{\mathbb{K}}(\mathbf{I}_1, \mathbf{q}) &\rightarrow k\{r_1, r_2, \theta\}, g, \sigma, \\ \mathbf{I}'_2 &= (g \cdot \mathbf{I}_1) \otimes k, \\ \mathbf{I}_2 &= \mathbf{I}'_2 + (g \cdot \mathbf{I}_1 - \mathbf{I}'_2) \cdot \sigma. \end{aligned} \quad (4)$$

White Balance & CCM Matrix: White balance (WB) mimics the color constancy of the human vision system (HVS) by aligning “white” color with the white object, resulting in the captured image reflecting the combination of light color and material reflectance [1,2,14]. In our work, we hope to find an adaptive white balance for different images under various lighting scenarios. Motivated by the design of Shades of Gray (SoG) [20] WB algorithm, where gray-world WB and Max-RGB WB can be regarded as a subcase, we’ve employed a learnable parameter ρ to replace the gray-world’s L-1 average with an adaptive Minkowski distance average (see Eq. 5), the QAL block $\mathbb{P}_{\mathbb{M}}$ predicts Minkowski distance’s hyper-parameter ρ , a demonstration of various ρ is shown in Fig. 4. After finding the suitable ρ , \mathbf{I}_2 would multiply to white balance matrix to generate \mathbf{I}_3 .

Color Conversion Matrix (CCM) within the ISP is constrained by the specific camera model. Here we standardize the CCM as a single learnable 3×3 matrix, initialized as a unit diagonal matrix \mathbf{E}_3 . The QAL block $\mathbb{P}_{\mathbb{M}}$ predicts 9 parameters here, which are then added to \mathbf{E}_3 to form the final 3×3 matrix \mathbf{E}_{ccm} . Then \mathbf{I}_3 would multiply to CCM \mathbf{E}_{ccm} to generate \mathbf{I}_4 , equation as follow:

$$\begin{aligned} \mathbb{P}_{\mathbb{M}}(\mathbf{I}_2, \mathbf{q}) &\rightarrow \rho, \mathbf{E}_{ccm}, \\ m_{i \in (r, g, b)} &= \sqrt[\rho]{\text{avg}(\mathbf{I}_2(i)^\rho)} / \sqrt[\rho]{\text{avg}((\mathbf{I}_2)^\rho)}, \\ \mathbf{I}_3 &= \mathbf{I}_2 * \begin{bmatrix} m_r & & \\ & m_g & \\ & & m_b \end{bmatrix}, \quad \mathbf{I}_4 = \mathbf{I}_3 * \mathbf{E}_{ccm}. \end{aligned} \quad (5)$$




Fig. 4: Left, we use query adaptive learning (QAL) to predict key parameters is ISP process. Right, we show RAW-Adapter different blocks' parameter.

Color Manipulation Process: In ISP, color manipulation process is commonly achieved through lookup table (LUT), such as 1D and 3D LUTs [16]. Here, we consolidate color manipulation operations into a single 3D LUT, adjusting the color of image I_4 to produce output image I_5 . Leveraging advancements in LUT techniques, we choose the latest neural implicit 3D LUT (NILUT) [11], for its speed efficiency and ability to facilitate end-to-end learning in our pipeline [3]. We denote NILUT [11] as \mathbb{L} , \mathbb{L} maps the input pixel intensities R, G, B to a continuous coordinate space, followed by the utilization of implicit neural representation (INR) [61], which involves using a multi-layer perceptron (MLP) network to map the output pixel intensities to R', G', B' :

$$\mathbf{I}_5(R', G', B') = \mathbb{L}(\mathbf{I}_4(R, G, B)). \quad (6)$$

Image \mathbf{I}_5 obtained through image-level adapters will be forwarded to the downstream network's backbone. Furthermore, \mathbf{I}_5 's features obtained in backbone will be fused with model-level adapters, which we will discuss in detail.

3.2 Model-level Adapters

Input-level adapters guarantee the production of machine vision-oriented image \mathbf{I}_5 for high-level models. However, information at the ISP stages ($\mathbf{I}_1 \sim \mathbf{I}_4$) is almost overlooked. Inspired by adapter design in current NLP and CV area [8, 28, 33], we employ the prior information from the ISP stage as model-level adapters to guide subsequent models' perception. Additionally, model-level adapters promote a tighter integration between downstream tasks and the ISP stage. Dotted lines in Fig. 3(a) represent model-level adapters.

Model-level adapters \mathbb{M} integrate the information from ISP stages to the network backbone. As shown in Fig. 3(a), we denote the stage 1 ~ 4 as different stages in network backbone, image \mathbf{I}_5 would pass by backbone stages 1 ~ 4 and then followed by detection or segmentation head. We utilize convolution layer c_i to extract features from \mathbf{I}_1 through \mathbf{I}_4 , these extracted features are concatenated as $\mathbb{C}(\mathbf{I}_{1\sim 4}) = \mathbb{C}(c_1(\mathbf{I}_1), c_2(\mathbf{I}_2), c_3(\mathbf{I}_3), c_4(\mathbf{I}_4))$. Subsequently, $\mathbb{C}(\mathbf{I}_{1\sim 4})$ go through

³ To save memory occupy, we change the channel number 128 in [11] to 32.

two residual blocks [25] and generate adapter f_1 , f_1 would merge with stage 1’s feature by a merge block, detail structure of merge block is shown in Fig. 3(c), which roughly includes concatenate process and a residual connect, finally merge block would output network feature for stage 2 and an additional adapter f_2 . Then process would repeat in stage 2 and stage 3 of network backbone. We collectively refer to all structures related to model-level adapters as \mathbb{M} .

We show RAW-Adapter’s different part parameter number in Fig. 4 right, including input-level adapters $\{\mathbb{P}_K (37.57K), \mathbb{P}_M (37.96K), \mathbb{L} (1.97K)\}$ and model level adapters $\mathbb{M} (114.87K \sim 687.59K)$, model level adapters’ parameter number depend on following network. Total parameter number of RAW-Adapter is around 0.2M to 0.8M, much smaller than following backbones ($\sim 25.6M$ in ResNet-50 [25], $\sim 197M$ in Swin-L [45]), also smaller than previous SOTA methods like SID [6] (11.99M) and Dirty-Pixel [18] (4.28M). Additionally, subsequent experiments will demonstrate that the performance of RAW-Adapter on backbone networks with fewer parameters is even superior to that of previous algorithms on backbone networks with higher parameters.

4 Experiments

4.1 Dataset and Experimental Setting

Dataset. We conducted experiments on object detection and semantic segmentation tasks, utilizing a combination of various synthetic and real-world RAW image datasets, an overview of the datasets can be found in Table 1.

For the object detection task, we take 2 open source real-world dataset PASCAL RAW [52] and LOD [27]. PASCAL RAW [52] is a normal-light condition dataset with 4259 RAW images, taken by a Nikon D3200 DSLR camera with 3 object classes. To verify the generalization capability of RAW-Adapter across various lighting conditions, we additionally synthesized low-light and overexposure datasets based on the PASCAL RAW dataset, named PASCAL RAW (dark) and PASCAL RAW (over-exp) respectively, the synthesized datasets are identical to the PASCAL RAW dataset in all aspects except for brightness levels. For PASCAL RAW (dark) and PASCAL RAW (over-exp) synthesis, the light intensity and environment irradiance exhibit a linear relationship on RAW images, thus we employed the synthesis method from the previous work [14, 78]:

$$\begin{aligned} x_n &\sim N(\mu = lx, \sigma^2 = \delta_r^2 + \delta_s lx) \\ y &= lx + x_n, \end{aligned} \tag{7}$$

where x denotes the original normal-light RAW image and y denotes degraded RAW image, $\delta_s = \sqrt{S}$ denotes shot noise while \sqrt{S} is signal of the sensor, δ_r denotes read noise, l denotes the light intensity parameter which randomly chosen from [0.05, 0.4] in PASCAL RAW (dark), and randomly chosen from [1.5, 2.5] in PASCAL RAW (over-exp). Additionally, we follow PASCAL RAW [52]’s dataset split to separate the training set and test set.

Table 1: Dataset and framework setting in our experiments.

	PASCAL RAW	PASCAL RAW (dark/over-exp)	LOD	ADE20K RAW	ADE20K RAW (dark/over-exp)
Task	Object Detection			Semantic Segmentation	
Number	4259		2230	27,574	
Type	real-world	synthesis	real-world	synthesis	
Sensor	Nikon D3200 DSLR		Canon EOS 5D Mark IV	-	
Framework	RetinaNet [43] & Sparse-RCNN [64]			Segformer [70]	
Backbone	ResNet [25]			MIT [70]	
pre-train	ImageNet [44] pre-train weights				

Meanwhile, LOD [27] is a real-world dataset with 2230 low-light condition RAW images, taken by a Canon EOS 5D Mark IV camera with 8 object classes, we take 1800 images as training set and the other 430 images as test set.

For the semantic segmentation task, we utilized the widely-used sRGB dataset ADE20K [80] to generate RAW dataset. Leveraging state-of-the-art unprocessing methods InvISP [71], we synthesized RAW images corresponding to the ADE20K sRGB dataset. Using InvISP, we projected input sRGB images into RAW format, effectively creating an ADE20K RAW dataset. To simulate various lighting conditions, we employed the same synthesis method from PASCAL RAW (dark/over-exp) to generate low-light and over-exposure RAW images, name as ADE20K RAW (dark/over-exp), training & test split is same as ADE20K.

Implement Details. We build our framework based on the open-source computer vision toolbox `mmdetection` [7] and `mmsegmentation` [12], both object detection tasks and semantic segmentation tasks are initialed with ImageNet pre-train weights (see Table. 1), and we apply the data augmentation pipeline in the default setting, mainly include random crop, random flip, and multi-scale test, etc. For the object detection task, we adopt the 2 mainstream object detectors: RetinaNet [43] and Sparse-RCNN [64] with ResNet [25] backbone. For the semantic segmentation task, we choose to use the mainstream segmentation framework Segformer [70] with MIT [70] backbone⁴.

Comparison Methods. We conducted comparative experiments with the current state-of-the-art (SOTA) algorithms, including various open-sourced ISP methods [6, 34, 35, 71, 79] and joint-training method DirtyPixels [18], among these Karaimer *et al.* [35] is a traditional ISP method with various human manipulate steps, and InvISP [71] & Lite-ISP [79] & SID [6] & DNF [34] are current SOTA network-based ISP methods, where SID [6] and DNF [34] is especially for the low-light condition RAW data. For fairness compared with the above ISP methods, both the training and test RAW data are rendered using the respective

⁴ For more experiments on other frameworks, please refer to supplementary materials.

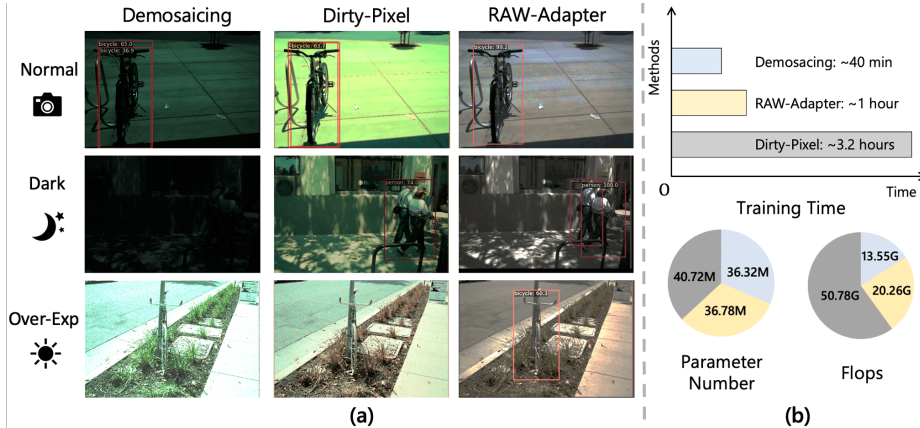


Fig. 5: (a). Detection performance on **PASCAL RAW** [52] (normal/dark/over-exp). (b). Efficiency comparison (blue: vanilla, yellow: RAW-Adapter, gray: Dirty-Pixel [18]).

compared ISP. DirtyPixels [18] is the current SOTA joint-training method which uses a stack of residual UNet [58] as a low-level processor and joint optimization with the following task-specific networks. For fairness, all comparison methods adopt the same data augmentation process and the same training setting, we will introduce detailed experiment results in the following section.

4.2 Object Detection Evaluation

For object detection task on PASCAL RAW [52] dataset, we adopt RetinaNet [43] with different size ResNet [25] backbones (ResNet-18, ResNet-50), all the models are trained on a single NVIDIA Tesla V100 GPU with SGD optimizer, the batch size is set to 4, training images are cropped into range of (400, 667) and training epochs are set to 50. Table. 2 shows the detection results with (a). ResNet-18 and (b). ResNet-50 backbone, with comparisons of demosaiced RAW data (“Demosaicing”), camera default ISP in [52], various current SOTA ISP solutions [6, 34, 35, 79] and Dirty-Pixel [18], we can observe that sometimes ISP algorithms in normal and over-exposure conditions may even degraded the detection performance, and ISP in dark scenarios sometimes could improves detection performance, additionally manual design ISP solution [35] even out-perform the deep learning solutions. Previous joint-training method Dirty-Pixel [18] can improve detection performance. Overall, our RAW-Adapter method achieves the best performance and even outperforms some ISP algorithms utilizing ResNet-50 backbone when employing ResNet-18 backbone, which can significantly reduce computational load and model parameters.

Visualization of PASCAL RAW’s detection results are shown in Fig. 5(a), we show the detection results under different lightness conditions, background are images generated from Dirty-Pixel [18] and RAW-Adapter (I₅). Our method could achieve satisfactory detection results across different lightness, while other

Table 2: Comparison results on **PASCAL RAW** dataset [52], we take [43] with (a). ResNet-18 and (b). ResNet-50 backbone, **bold** denotes the best result.

(a) ResNet-18.				(b) ResNet-50.			
Method	mAP			Method	mAP		
	normal	over-exp	dark		normal	over-exp	dark
Default ISP	88.3	-	-	Default ISP	89.4	-	-
Demosacing	87.7	87.7	80.3	Demosacing	89.2	88.8	82.6
Karaimer <i>et al.</i> [35]	86.0	85.6	81.9	Karaimer <i>et al.</i> [35]	89.4	86.8	84.6
Lite-ISP [79]	85.2	84.2	71.9	Lite-ISP [79]	89.3	85.1	73.5
InvISP [71]	84.1	86.6	70.9	InvISP [71]	89.1	87.3	74.7
SID [6]	-	-	78.2	SID [6]	-	-	81.5
DNF [34]	-	-	81.1	DNF [34]	-	-	82.8
Dirty-Pixel	88.6	88.0	80.8	Dirty-Pixel	89.0	89.0	83.6
RAW-Adapter	88.7	88.7	82.5	RAW-Adapter	89.6	89.4	86.6

Table 3: Comparison with ISP methods and Dirty-Pixel on **LOD** dataset [27]. We show the detection performance (mAP) \uparrow of RetinaNet (R-Net) [43] and Sparse-RCNN (Sp-RCNN) [64], **bold** denotes the best result.

Methods	Demosacing	Karaimer <i>et al.</i> [35]	InvISP [71]	SID [6]
mAP(R-Net)	58.5	54.4	56.9	49.1
mAP(Sp-RCNN)	57.7	52.2	49.4	43.1
Methods	DNF [34]	Dirty-Pixel [18]	RAW-Adapter (w/o M)	RAW-Adapter
mAP(R-Net)	58.8	61.6	61.6	62.2
mAP(Sp-RCNN)	55.8	58.8	58.6	59.2

methods face false alarm and miss detect. Efficiency comparison with Dirty-Pixel is shown in Fig. 5(b), our algorithm has achieved significant improvements in both training time acceleration and savings in Flops/Parameter Number.

Detection performance on LOD [27] is shown in Table. 3, we adopt two detector RetinaNet [43] and Sparse-RCNN [64], both with ResNet-50 [25] backbone, training epochs are set to 35, Sparse-RCNN [64] are trained by 4 GPUs with Adam optimizer. With comparing of various ISP methods [6, 34, 35, 79], Dirty-Pixel [18] and RAW-Adapter with only input-level adapters (w/o M). Table. 3 show that our algorithm can achieve the best performance on both two detectors.

4.3 Semantic Segmentation Evaluation

For semantic segmentation on ADE20K RAW dataset, we choose Segformer [70] as the segmentation framework with the different size MIT [70] backbones (MIT-B5, MIT-B3, MIT-B0), all the models are trained on 4 NVIDIA Tesla V100 GPU

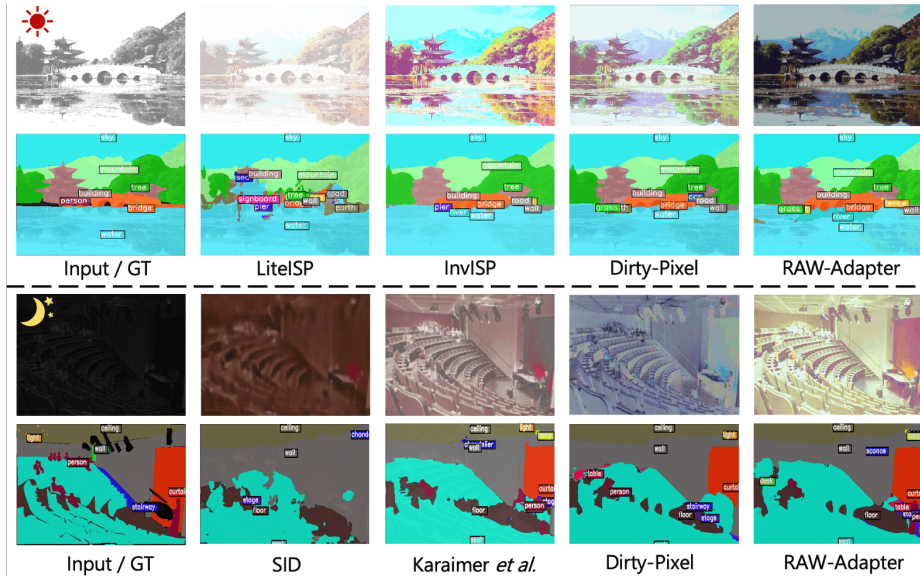


Fig. 6: Semantic segmentation results on over-exposure and low-light RAW data, compare with various ISP methods [6, 35, 71, 79] and Dirty-Pixel [18].

with Adam optimizer, the batch size is set to 4, training images are cropped into 512×512 and training iters are set to 80000. We make comparison with single demosaicing, various ISP methods [6, 34, 35, 71, 79] and Dirty-Pixel [18], and RAW-Adapter with only input-level adapters (w/o \mathcal{M}).

Comparison results are shown in Table. 4, where we compare both efficiency (parameters \downarrow , inference time \downarrow) and performance (mIOU \uparrow), inference time is calculated on a single Tesla V100 GPU. We first show the results on vanilla demosaiced RAW data, then we show the results on RAW data processed by various ISP methods [6, 34, 35, 71, 79], followed by results of Dirty-Pixel [18], and our method without and with model-level adapter \mathcal{M} . For Table. 4, we can find that directly employing ISP methods sometimes fails to enhance the performance of downstream segmentation tasks, particularly in low-light conditions. Joint-training method Dirty-Pixel [18] could improve down-stream performance and sometimes even degrade segmentation performance, but increase more network parameters. Overall, our proposed RAW-Adapter could effectively enhance performance while adding a slight parameter count and improving inference efficiency. RAW-Adapter on backbones with fewer parameters can even outperform previous methods on backbones with more parameters (*i.e.* The RAW-Adapter, employing the MIT-B3 backbone, achieves a dark condition mIOU of 40.52, surpassing other methods utilizing the MIT-B5 backbone despite having fewer than $\sim 36\text{M}$ parameters.).

The visualization of segmentation results are shown in Fig. 6, where row 1 and row 3 depict images rendered by different ISPs [6, 35, 71, 79] and images

Table 4: Comparison with previous methods on **ADE 20K RAW** (normal/over-exp/dark). **Bold** denotes the best result while underline denotes second best result.

	backbone	params(M) ↓	inference time(s) ↓	mIOU ↑ (normal)	mIOU ↑ (over-exp)	mIOU ↑ (dark)
Demosacing	MIT-B5	82.01	0.105	45.26	42.08	30.97
Karaimer <i>et al.</i> [35]			0.525	46.42	43.71	39.79
InvISP [71]			0.203	47.53	43.55	4.87
LiteISP [79]			0.261	43.22	42.01	5.52
DNF [34]			0.186	-	-	25.82
SID [6]			0.312	-	-	26.44
Dirty-Pixel [18]	MIT-B3	86.29	0.159	47.76	<u>46.80</u>	40.02
	MIT-B0	48.92	0.098	47.19	45.33	38.93
	MIT-B0	8.00	<u>0.049</u>	35.43	32.10	27.19
RAW-Adapter (w/o M)	MIT-B5	82.09	0.148	<u>48.33</u>	46.48	<u>40.91</u>
	MIT-B3	44.72	0.086	47.01	45.00	39.34
	MIT-B0	3.80	0.032	35.09	32.82	27.29
RAW-Adapter	MIT-B5	82.53	0.167	48.38	47.06	41.82
	MIT-B3	45.16	0.102	47.59	45.66	40.52
	MIT-B0	<u>4.26</u>	0.053	35.67	33.44	27.78

generated from Dirty-Pixel [18] and RAW-Adapter (I_5), row 2 and row 4 denote our segmentation results appear superior compared to other methods. We were even pleased to discover that, although we do not add any human-vision-oriented loss function to constrain the model, RAW-Adapter still produces visually satisfactory image I_5 (see Fig. 6). This may be attributed to our preservation of traditional ISP stages in input-level adapters. Due to page limitations, we refer to more experimental results, visualization results, and ablation analysis of our method in the supplementary part.

5 Conclusion

In this paper, we introduce RAW-Adapter, an effective solution for adapting pre-trained sRGB models to camera RAW data. With input-level adapters and model-level adapters working in tandem, RAW-Adapter effectively forwards RAW images from various lighting conditions to downstream vision tasks, our method has achieved state-of-the-art (SOTA) results across multiple datasets.

For future research directions, we believe it is feasible to design a unified model capable of adapting to RAW visual tasks under different lighting conditions, without the need for retraining in each lighting scenario like RAW-Adapter. Additionally, designing multi-task decoders can enable the accommodation of various tasks, facilitating more effective integration of different visual tasks on RAW images with large-scale models.

Acknowledgements

This research is partially supported by JST Moonshot R&D Grant Number JPMJPS2011, CREST Grant Number JPMJCR2015 and Basic Research Grant (Super AI) of Institute for AI and Beyond of the University of Tokyo.

References

1. Affi, M., Brown, M.S.: What else can fool deep learning? addressing color constancy errors on deep neural network performance. 2019 IEEE/CVF International Conference on Computer Vision (ICCV) pp. 243–252 (2019), <https://api.semanticscholar.org/CorpusID:207995696>
2. Brooks, T., Mildenhall, B., Xue, T., Chen, J., Sharlet, D., Barron, J.T.: Unprocessing images for learned raw denoising. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11036–11045 (2019)
3. Buckler, M., Jayasuriya, S., Sampson, A.: Reconfiguring the imaging pipeline for computer vision. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 975–984 (2017)
4. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European conference on computer vision. pp. 213–229. Springer (2020)
5. Chatterjee, P., Joshi, N., Kang, S.B., Matsushita, Y.: Noise suppression in low-light images through joint denoising and demosaicing. In: CVPR 2011. pp. 321–328. IEEE (2011)
6. Chen, C., Chen, Q., Xu, J., Koltun, V.: Learning to see in the dark. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3291–3300 (2018)
7. Chen, K., Wang, J., et al.: Mmdetection: Open mmlab detection toolbox and benchmark (2019)
8. Chen, Z., Duan, Y., Wang, W., He, J., Lu, T., Dai, J., Qiao, Y.: Vision transformer adapter for dense predictions. In: The Eleventh International Conference on Learning Representations (2023), <https://openreview.net/forum?id=plKu2GByCNW>
9. Cheng, B., Schwing, A.G., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. In: NeurIPS (2021)
10. Conde, M.V., McDonagh, S., Maggioni, M., Leonardi, A., Pérez-Pellitero, E.: Model-based image signal processors via learnable dictionaries. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 481–489 (2022)
11. Conde, M.V., Vazquez-Corral, J., Brown, M.S., Timofte, R.: Nilut: Conditional neural implicit 3d lookup tables for image enhancement. In: Proceedings of the AAAI Conference on Artificial Intelligence (2024)
12. Contributors, M.: MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/msegmentation> (2020)
13. Cui, Z., Li, K., Gu, L., Su, S., Gao, P., Jiang, Z., Qiao, Y., Harada, T.: You only need 90k parameters to adapt light: a light weight transformer for image enhancement and exposure correction. In: BMVC. p. 238 (2022)
14. Cui, Z., Qi, G.J., Gu, L., You, S., Zhang, Z., Harada, T.: Multitask aet with orthogonal tangent regularity for dark object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2553–2562 (October 2021)

15. Cui, Z., Zhu, Y., Gu, L., Qi, G.J., Li, X., Zhang, R., Zhang, Z., Harada, T.: Exploring resolution and degradation clues as self-supervised signal for low quality object detection. In: European Conference on Computer Vision. pp. 473–491. Springer (2022)
16. Delbracio, M., Kelly, D., Brown, M.S., Milanfar, P.: Mobile computational photography: A tour. *Annual Review of Vision Science* **7**(1), 571–604 (2021). <https://doi.org/10.1146/annurev-vision-093019-115521>, <https://doi.org/10.1146/annurev-vision-093019-115521>, PMID: 34524880
17. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A Large-Scale Hierarchical Image Database. In: CVPR09 (2009)
18. Diamond, S., Sitzmann, V., Julca-Aguilar, F., Boyd, S., Wetzstein, G., Heide, F.: Dirty pixels: Towards end-to-end image processing and perception. *ACM Transactions on Graphics (SIGGRAPH)* (2021)
19. Dong, X., Yokoya, N.: Understanding dark scenes by contrasting multi-modal observations. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 840–850 (January 2024)
20. Finlayson, G.D., Trezzi, E.: Shades of gray and colour constancy. In: Color and Imaging Conference. vol. 2004, pp. 37–41. Society for Imaging Science and Technology (2004)
21. Goyal, B., Lalonde, J.F., Li, Y., Gupta, M.: Robust scene inference under noise-blur dual corruptions. 2022 IEEE International Conference on Computational Photography (ICCP) pp. 1–12 (2022), <https://api.semanticscholar.org/CorpusID:251040807>
22. Guo, Y., Luo, F., Wu, X.: Learning degradation-independent representations for camera isp pipelines. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 25774–25783 (June 2024)
23. Hansen, N., Ostermeier, A.: Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. In: Proceedings of IEEE International Conference on Evolutionary Computation. pp. 312–317 (1996). <https://doi.org/10.1109/ICEC.1996.542381>
24. Hasinoff, S.W., Sharlet, D., Geiss, R., Adams, A., Barron, J.T., Kainz, F., Chen, J., Levoy, M.: Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* **35**(6) (2016)
25. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
26. Heide, F., Steinberger, M., Tsai, Y.T., Rouf, M., Pająk, D., Reddy, D., Gallo, O., Liu, J., Heidrich, W., Egiazarian, K., et al.: Flexisp: A flexible camera image processing framework. *ACM Transactions on Graphics (ToG)* **33**(6), 1–13 (2014)
27. Hong, Y., Wei, K., Chen, L., Fu, Y.: Crafting object detection in very low light. In: BMVC (2021)
28. Houlsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., Gelly, S.: Parameter-efficient transfer learning for nlp. In: International Conference on Machine Learning. pp. 2790–2799. PMLR (2019)
29. Hu, Y., He, H., Xu, C., Wang, B., Lin, S.: Exposure: A white-box photo post-processing framework. *ACM Transactions on Graphics (TOG)* **37**(2), 26 (2018)
30. Ignatov, A., Van Gool, L., Timofte, R.: Replacing mobile camera isp with a single deep learning model. arXiv preprint arXiv:2002.05509 (2020)

31. Iscen, A., Zhang, J., Lazebnik, S., Schmid, C.: Memory-efficient incremental learning through feature adaptation. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16. pp. 699–715. Springer (2020)
32. Jensen, H.W., Durand, F., Dorsey, J., Stark, M.M., Shirley, P., Premože, S.: A physically-based night sky model. In: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques. p. 399–408. SIGGRAPH '01, Association for Computing Machinery, New York, NY, USA (2001). <https://doi.org/10.1145/383259.383306>, <https://doi.org/10.1145/383259.383306>
33. Jia, M., Tang, L., Chen, B.C., Cardie, C., Belongie, S., Hariharan, B., Lim, S.N.: Visual prompt tuning. In: European Conference on Computer Vision. pp. 709–727. Springer (2022)
34. Jin, X., Han, L., Li, Z., Chai, Z., Guo, C., Li, C.: Dnf: Decouple and feedback network for seeing in the dark. In: ICCV (2023)
35. Karaimer, H.C., Brown, M.S.: A software platform for manipulating the camera imaging pipeline. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) Computer Vision – ECCV 2016. pp. 429–444. Springer International Publishing, Cham (2016)
36. Ke, L., Ye, M., Danelljan, M., Liu, Y., Tai, Y.W., Tang, C.K., Yu, F.: Segment anything in high quality. In: NeurIPS (2023)
37. Kim, W., Kim, G., Lee, J., Lee, S., Baek, S.H., Cho, S.: Paramisp: Learned forward and inverse isps using camera parameters. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 26067–26076 (2024)
38. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 4015–4026 (October 2023)
39. Koskinen, S., Yang, D., Kämäräinen, J.K.: Reverse imaging pipeline for raw rgb image augmentation. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 2896–2900 (2019). <https://doi.org/10.1109/ICIP.2019.8804406>
40. Li, Z., Lu, M., Zhang, X., Feng, X., Asif, M.S., Ma, Z.: Efficient visual computing with camera raw snapshots. TPAMI pp. 1–18 (2024). <https://doi.org/10.1109/TPAMI.2024.3359326>
41. Li, Z., Jiang, H., Cao, M., Zheng, Y.: Polarized color image denoising. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9873–9882 (2023). <https://doi.org/10.1109/CVPR52729.2023.00952>
42. Liang, Z., Cai, J., Cao, Z., Zhang, L.: Cameranet: A two-stage framework for effective camera isp learning. IEEE Transactions on Image Processing **30**, 2248–2262 (2021). <https://doi.org/10.1109/TIP.2021.3051486>
43. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
44. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) Computer Vision – ECCV 2014. pp. 740–755. Springer International Publishing, Cham (2014)
45. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
46. Mildenhall, B., Hedman, P., Martin-Brualla, R., Srinivasan, P.P., Barron, J.T.: Nerf in the dark: High dynamic range view synthesis from noisy raw images. In:

- Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16190–16199 (2022)
47. Monakhova, K., Richter, S.R., Waller, L., Koltun, V.: Dancing under the stars: Video denoising in starlight. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16241–16251 (June 2022)
 48. Mosleh, A., Sharma, A., Onzon, E., Mannan, F., Robidoux, N., Heide, F.: Hardware-in-the-loop end-to-end optimization of camera image processing pipelines. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
 49. Nam, S., Kim, S.J.: Modelling the scene dependent imaging in cameras with a deep neural network. 2017 IEEE International Conference on Computer Vision (ICCV) pp. 1726–1734 (2017), <https://api.semanticscholar.org/CorpusID:23920472>
 50. Nguyen, R.M.H., Brown, M.S.: Raw image reconstruction using a self-contained srgb-jpeg image with only 64 kb overhead. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1655–1663 (2016). <https://doi.org/10.1109/CVPR.2016.183>
 51. Nishimura, J., Gerasimow, T., Sushma, R., Sutic, A., Wu, C.T., Michael, G.: Automatic isp image quality tuning using nonlinear optimization. In: 2018 25th IEEE International Conference on Image Processing (ICIP). pp. 2471–2475. IEEE (2018)
 52. Omid-Zohoor, A., Ta, D., Murmann, B.: Pascalraw: raw image database for object detection (2014)
 53. Onzon, E., Mannan, F., Heide, F.: Neural auto-exposure for high-dynamic range object detection. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7706–7716 (2021). <https://doi.org/10.1109/CVPR46437.2021.00762>
 54. Potlapalli, V., Zamir, S.W., Khan, S., Khan, F.: Promptir: Prompting for all-in-one image restoration. In: Thirty-seventh Conference on Neural Information Processing Systems (2023)
 55. Qin, H., Han, L., Wang, J., Zhang, C., Li, Y., Li, B., Hu, W.: Attention-aware learning for hyperparameter prediction in image processing pipelines. In: European Conference on Computer Vision. pp. 271–287. Springer (2022)
 56. Ramanath, R., Snyder, W.E., Yoo, Y., Drew, M.S.: Color image processing pipeline. IEEE Signal processing magazine **22**(1), 34–43 (2005)
 57. Rebuffi, S.A., Bilen, H., Vedaldi, A.: Learning multiple visual domains with residual adapters. Advances in neural information processing systems **30** (2017)
 58. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. pp. 234–241. Springer International Publishing, Cham (2015)
 59. Sayed, M., Brostow, G.: Improved handling of motion blur in online object detection. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1706–1716 (2021). <https://doi.org/10.1109/CVPR46437.2021.00175>
 60. Schwartz, E., Giryes, R., Bronstein, A.M.: Deepisp: Toward learning an end-to-end image processing pipeline. IEEE Transactions on Image Processing **28**(2), 912–923 (2018)
 61. Sitzmann, V., Martel, J.N., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: Proc. NeurIPS (2020)
 62. Souza, M., Heidrich, W.: MetaISP – Exploiting Global Scene Structure for Accurate Multi-Device Color Rendition. In: Guthe, M., Grosch, T. (eds.) Vision, Modeling,

- and Visualization. The Eurographics Association (2023). <https://doi.org/10.2312/vmv.20231236>
63. Stickland, A.C., Murray, I.: Bert and pals: Projected attention layers for efficient adaptation in multi-task learning. In: International Conference on Machine Learning. pp. 5986–5995. PMLR (2019)
 64. Sun, P., Zhang, R., Jiang, Y., Kong, T., Xu, C., Zhan, W., Tomizuka, M., Li, L., Yuan, Z., Wang, C., Luo, P.: SparseR-CNN: End-to-end object detection with learnable proposals. arXiv preprint arXiv:2011.12450 (2021)
 65. Sung, Y.L., Cho, J., Bansal, M.: Vl-adapter: Parameter-efficient transfer learning for vision-and-language tasks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5227–5237 (2022)
 66. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
 67. Wang, Y., Yu, Y., Yang, W., Guo, L., Chau, L.P., Kot, A.C., Wen, B.: Raw image reconstruction with learned compact metadata. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 18206–18215 (June 2023)
 68. Wei, K., Fu, Y., Yang, J., Huang, H.: A physics-based noise formation model for extreme low-light raw denoising. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
 69. Wu, C.T., Isikdogan, L.F., Rao, S., Nayak, B., Gerasimow, T., Sutic, A., Ainkedem, L., Michael, G.: Visionisp: Repurposing the image signal processor for computer vision applications. In: 2019 IEEE International Conference on Image Processing (ICIP). IEEE (Sep 2019). <https://doi.org/10.1109/icip.2019.8803607>, <http://dx.doi.org/10.1109/ICIP.2019.8803607>
 70. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. arXiv preprint arXiv:2105.15203 (2021)
 71. Xing, Y., Qian, Z., Chen, Q.: Invertible image signal processing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6287–6296 (2021)
 72. Xu, R., Chen, C., Peng, J., Li, C., Huang, Y., Song, F., Yan, Y., Xiong, Z.: Toward raw object detection: A new benchmark and a new model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13384–13393 (2023)
 73. Yoshimura, M., Otsuka, J., Irie, A., Ohashi, T.: Dynamicisp: Dynamically controlled image signal processor for image recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 12866–12876 (October 2023)
 74. Yoshimura, M., Otsuka, J., Irie, A., Ohashi, T.: Rawgment: Noise-accounted raw augmentation enables recognition in a wide variety of environments. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14007–14017 (June 2023)
 75. Yu, K., Li, Z., Peng, Y., Loy, C.C., Gu, J.: Reconfigisp: Reconfigurable camera image processing pipeline. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4248–4257 (2021)
 76. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Cycleisp: Real image restoration via improved data synthesis. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2696–2705 (2020)

77. Zhang, R., Zhang, W., Fang, R., Gao, P., Li, K., Dai, J., Qiao, Y., Li, H.: Tip-adapter: Training-free adaption of clip for few-shot classification. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *Computer Vision – ECCV 2022*. pp. 493–510. Springer Nature Switzerland, Cham (2022)
78. Zhang, Y., Zhang, Y., Zhang, Z., Zhang, M., Tian, R., Ding, M.: Isp-teacher: Image signal process with disentanglement regularization for unsupervised domain adaptive dark object detection. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 38, pp. 7387–7395 (2024)
79. Zhang, Z., Wang, H., Liu, M., Wang, R., Zuo, W., Zhang, J.: Learning raw-to-srgb mappings with inaccurately aligned supervision. In: *ICCV* (2021)
80. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 633–641 (2017)
81. Zhou, W., Gao, S., Zhang, L., Lou, X.: Histogram of oriented gradients feature extraction from raw bayer pattern images. *IEEE Transactions on Circuits and Systems II: Express Briefs* **67**(5), 946–950 (2020). <https://doi.org/10.1109/TCSII.2020.2980557>