

DynamicISP: Dynamically Controlled Image Signal Processor for Image Recognition

Masakazu Yoshimura Junji Otsuka Atsushi Irie Takeshi Ohashi
Sony Group Corporation

{masakazu.yoshimura, junji.otsuka, atsushi.irie, takeshi.a.ohashi}@sony.com

Abstract

Image Signal Processors (ISPs) play important roles in image recognition tasks as well as in the perceptual quality of captured images. In most cases, experts make a lot of effort to manually tune many parameters of ISPs, but the parameters are sub-optimal. In the literature, two types of techniques have been actively studied: a machine learning-based parameter tuning technique and a DNN-based ISP technique. The former is lightweight but lacks expressive power. The latter has expressive power, but the computational cost is too heavy on edge devices. To solve these problems, we propose “DynamicISP,” which consists of multiple classical ISP functions and dynamically controls the parameters of each frame according to the recognition result of the previous frame. We show our method successfully controls the parameters of multiple ISP functions and achieves state-of-the-art accuracy with low computational cost in single and multi-category object detection tasks.

1. Introduction

The image signal processor (ISP) is an important component of modern digital cameras. It converts raw outputs of the sensors, RAW images, into commonly used standard RGB (sRGB) images, and is composed of many functions because of its multiple roles. For example, it calibrates the color sensitivity of each sensor with a color correction matrix, and improves image quality with demosaicing, denoising, and sharpening. Auto-exposure, auto-white balance, and tone mapping are necessary to mimic the adaptive and responsive characteristics of the human eye to produce images that closely resemble the appearance perceived by humans. The human eye cancels out the environment’s luminance with dark or light adaptation and also cancels out the environment’s color with color adaptation. As to the response characteristics, it is said that the human eye non-linearly responds to light intensity roughly obeying $y = x^{1/3}$ [38]. These functions for human eye adaptation




Figure 1. The concept of proposed DynamicISP. The parameters of the ISP are controlled dynamically based on what the downstream recognition model felt.

have another role. Optimizing these functions for compression enables conversion from high dynamic range signals in the real world to 8-bit sRGB images with little information loss. The real world has the range from 0.0001 [cd/m^2] (starlight) to 1.6 billion [cd/m^2] [34] (direct sunlight). We don’t go any further, but ISPs have more functions, like lens shading correction, dehazer, and bad pixel correction.

ISPs are also needed for image recognition tasks. DNN-based image recognition, which has achieved excellent results recently, shows even better performance with ISPs [10, 3]. However, ISP hyperparameters commonly require manual tuning by experts, which consumes considerable time [27]. The tuning is necessary for each sensor. It is ineffective and might be sub-optimal, especially for image recognition tasks. Easy-to-see images for humans are not equal to easy-to-recognize images. In fact, recent works demonstrate that the hyperparameter tuning for image recognition improves the accuracy [27, 46, 35, 39].

To achieve a better ISP for image recognition, we believe it is important to improve the expressive power of the ISP since most ISP functions are classical static functions. One solution is a DNN encoder-decoder-based ISP; many works indeed achieved state-of-the-art accuracy in image recognition performance [8, 26] and perceptual quality [33, 20]. However, the computation costs are enormous, especially when high-resolution outputs were required. Devices such as smartphones have limited computational resources and require an efficient ISP.

Based on the above, we propose “DynamicISP,” which consists of multiple conventional ISP functions, but their parameters are controlled dynamically per image. Although they are classical functions, their dynamic control compensates for their expressive power. Our controller is based on what the downstream recognition model felt and desired at the previous time step as shown in Fig. 1. Specifically, our method predicts appropriate ISP parameters for the following frame based on intermediate features of the downstream recognition model.

In terms of dynamic control, dynamic neural networks have been actively studied [44, 6] and a similar method, named NeuralAE, that controls auto-exposure with DNNs has recently been published [30]. Our method successfully controls the parameters of multiple ISP functions, which is difficult even with these successful dynamic control methods. In addition, despite its low computational cost, our method achieves state-of-the-art accuracy compared to encoder-decoder based DNN ISPs.

Our contributions are as follows:

- A novel ISP for image recognition that utilizes classical functions but dynamically controls their parameters based on what the downstream image recognition model senses.
- A residual output format of parameters, which decomposes the control into static tuning to average good parameters for the entire data set and further pursuit of the best parameters per image. It eases the difficulty of controlling complex functions.
- A latent update style ISP controller that manages multiple ISP functions. It generates shared latent variables which indicate how to change the RAW images with the entire ISP pipeline; then, after applying each ISP function, sequentially updates them to contain information on what the remaining functions should do considering upstream functions.
- End-to-end optimization method from ISP to image recognition model. A proposed *parameter initializer* module enabled the control of complex functions.
- Diverse evaluations show the effectiveness of the proposed DynamicISP. It achieves state-of-the-art accuracy despite the low computational cost.

2. Related Works

2.1. Image Signal Processor

In recent years, many ISPs that utilize machine learning have been proposed. They can be broadly classified into two categories: methods that tune the classical ISP functions

and methods that replace them with a DNN-based encoder-decoder model.

In the first category, they tune the existing ISP parameters using machine learning techniques. It defines perceptual loss or image recognition loss, depending on the goal, to increase perceptual quality or recognition accuracy. In some cases, evolutionary algorithms [32, 27, 13] or other search methods [28] are applied because most existing ISPs are black-box and the search space of the parameters have multiple local minima [39]. However, these methods optimize only ISP parameters for a fixed downstream part, and the performance gain is limited. To partly address the problem, a work proposes a joint training schedule wherein ISP and downstream CNN are optimized alternately while the other is fixed [35]. Other methods optimize the parameters with gradient-based backpropagation [43, 30, 45]. Although we mentioned NeuralAE [30] controls auto-exposure with DNN, ISP parameters are trained with backpropagation as static parameters. ReconfigISP [46] not only optimizes the hyperparameters but also explores functions by a neural architecture search method [19]. It is important to note that the computational costs of this category are low because their ISPs for inference time consist of classical functions. We argue that the problem with these methods is the lack of expressive power compared with the other category.

In the other category, methods create an ISP with encoder-decoder-based DNN. Recently, tremendous model architectures are proposed as an ISP function replacement, such as denoiser [48, 25], color constancy [29, 1], and deblurring [47, 42]. Some works propose generic models that can replace various functions with the same architecture [40, 5], while several others replace the whole ISP pipeline with DNN [33, 26, 8, 20]. Though DNN-based ISPs and ISP functions achieve state-of-the-art, the computational cost remains an issue especially when high-resolution outputs are necessary. Currently, smartphones are one of the major use cases of ISPs, thus low computational cost ISPs are highly desired.

Based on the above, we propose a dynamic ISP control method that doesn’t match either of these two categories. By dynamically controlling parameters, classical ISPs are empowered with a minimal computational cost. Our method is possible to optimize from ISP to image recognition end-to-end, going beyond the usual individual optimization of the ISP and the recognition model.

2.2. Low-Light Image Enhancement

It is said that conventional ISPs tend to be poor in a low-light environment [9, 15, 49]. Many works focus on improving perceptual quality or recognition accuracy in low-light environments. Most of them try to enhance the brightness of pre-captured sRGB images [23, 49, 15, 41, 9, 24]. These are

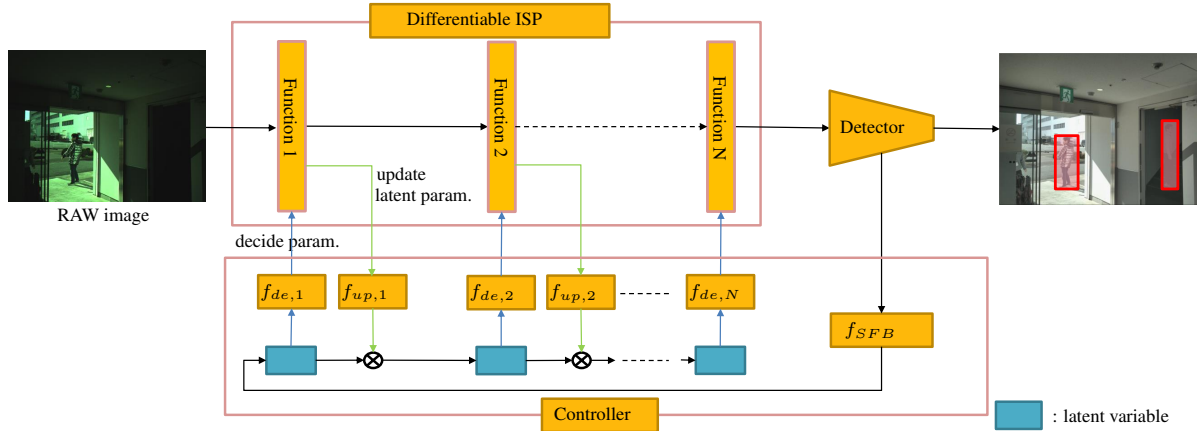


Figure 2. The proposed latent update style controller to control multi-layers.

very useful when improving already saved sRGB images but some information is already vanished due to the lossy conversion from RAW to sRGB. Other works optimize ISP for dark environment [33, 26, 20, 8]. Several works train recognition models to recognize RAW images, which have the richest information [14, 37].

Our first objective is to establish an ISP that can be used in any environment covering from low light to dazzling, high dynamic range (HDR), and even blurry. However, our DynamicISP can be useful when targeting only dark scenes because optimal ISP parameters can be different per environment even if the luminance is the same. Hence, we also evaluate our method on low-light image recognition.

2.3. Dynamic Parameter Control

The dynamic model is one of the current hot research topics. Attention-based feedforward control is successful in dynamically controlling DNN parameters. Several works control the amplification values of the convolution weights or PReLU [11] using a separate attention module fed features just before the convolution or PReLU [44, 6, 7]. These works prove the effectiveness of controlling partial DNN parameters.

However, these attention-based feedforward controls are difficult to apply to the parameters of the preprocessing part - there are no rich features before the target functions. NeuralAE [30] solves these situations by introducing a feedback circuit in which an intermediate feature of the detector’s backbone from the previous time step is used to control the exposure.

Although our ISP control can be based on feedback control as NeuralAE does, the control of multiple ISP functions with a controller has never been realized to the best of our knowledge; it fails when we simply apply the existing NeuralAE’s method. Therefore, we propose the feedback control method to control parameters of multiple ISP functions.

3. Methodology

In this section, the proposed dynamic ISP control and its training method are presented.

3.1. Dynamic ISP Control

Our problem setting for ISP control is to predict the next frame ($t + 1$) parameters $P = \{P_l | l = 1, \dots, L\}$ of ISP functions $I(x) = \{I_l(x) | l = 1, \dots, L\}$, where L is the number of ISP functions, to boost visibility for machine vision. Each function $I_l(x)$ contains N_l parameters $P_l = \{p_{l,n}^{t+1} | n = 1, \dots, N_l\}$. In a static frame scenario, we let the next frame ($t + 1$) just the same image as the previous frame t . The parameters P are estimated based on encoded intermediate feature X_t of the downstream recognition model in a feedback control format according to NeuralAE [30]. Specifically, we use the “*Semantic Feature Branch*” $f_{SFB}(X_t)$ proposed in NeuralAE to encode the features. While NeuralAE adds another histogram-based RAW image encoder independent from the recognition model, we did not use it because it was the bottleneck of the speed in our implementation, and no accuracy improvement was noticed, particularly when a large downstream recognition model was used.

Residual Output Format of Parameters

NeuralAE predicts the desired exposure by

$$p_{0,0}^{t+1} = f_{act}(f_{full}(f_{SFB}(X_t))), \quad (1)$$

where f_{full} and f_{act} are fully connected layers and an output activation layer respectively. Since ISP functions are not as simple as the exposure function, which just multiplies the input, we need to propose an elaborate output format to control them properly. We introduce learnable parameter $\hat{p}_{l,n}$ as

$$p_{l,n}^{t+1} = f_{act}(\hat{p}_{l,n} + f_{full}(f_{SFB}(X_t))). \quad (2)$$

By this formulation, the $p_{l,n}$ can be trained as an average good value for the entire data set, while the controller only predicts the gap between the average good value and the best value per image. It eases the difficulty of controlling complex functions. In detail, we define f_{act} as

$$f_{act}(x) = (p_{l,n,max} - p_{l,n,min}) \cdot \text{sigmoid}(x) + p_{l,n,min}, \quad (3)$$

to search in $(p_{l,n,min}, p_{l,n,max})$. This formulation is handy because we can easily set the search space of any hyper-parameters based on their domain. Some ISP parameters are limited in the range they can take. In the experiments, wide search ranges were set to allow free control within the ranges they can take.

Latent Update Style Controller

Both attention-based dynamic DNN [44, 6] and NeuralAE [30] predict parameters of only one function by a controller. However, multi-layer ISP control is necessary to predict the parameters of multi-layers. As modifying the parameters in upstream functions affect the following functions, controlling all functions is challenging. In fact, our experiments show that existing methods are difficult to control multi-layers. We tackle this issue with a proposed latent updated style controller. The proposal is illustrated in Fig. 2. First, the main controller f_{SFB} extracts what the downstream image recognition model felt and reacted to the previous frame and converts the information into a latent variable that contains what to do in overall functions. Then, after applying each function, the latent variable is updated to contain information on what to do in the remaining functions. The concept is realized with the following formula:

$$\begin{cases} P_l = f_{de,l}(V_{l-1}) = f_{act,l}(\hat{P}_l + f_{full,l}(V_{l-1})) \\ V_l = f_{up,l}(P_l, V_{l-1}) \end{cases}, \quad (4)$$

where V_{l-1} is the latent variable for deciding the parameters P_l of function I_l with $f_{de,l}$, and $f_{ud,l}$ updates it according to what conversion is done in I_l . In our experiments, we define $f_{ud,l}$ simply as

$$f_{up,l}(P_l, V_{l-1}) = f_a(P_l) \cdot V_{l-1}, \quad (5)$$

where $f_a(P_l)$ consists of a normalization layer, two fully connected layers, a ReLU activation between the fully connected layers, then an activation, $y = 5 \cdot \text{sigmoid}(x)$. The normalization layer is based on the range of the parameter search spaces: $(p_{l,n}^{t+1} - p_{l,n,min}) / (p_{l,n,max} - p_{l,n,min})$ for each n .

The latent update style disentangles the multi-layer control problem setting without adding image encoders between the functions to convert it into a single-layer control per each estimator problem. Adding image encoders per

function increases the computational cost compared to our 1-D latent variable manipulations.

3.2. Training Method

We mainly follow the two-frame training method of NeuralAE [30]. The first frame is input to the differentially implemented ISP and to a part of the image recognition backbone to get the intermediate feature X_t for ISP control. Second, the ISP is controlled for the next frame. Third, the next frame is input to the ISP and the image recognition model to compute a recognition loss. Finally, the recognition loss for the second frame is backpropagated through all paths above to train the ISP, controller, and recognition model end-to-end.

In NeuralAE, the exposure of the first frame is set to default, in other words, the exposure of the first frame is set to the same value used in data capturing. One major problem with ISP control is that it consists of non-linear functions and drastically changes the original RAW input depending on the parameters. To train the controller, it is fundamental to properly manage the ISP parameters for the first frame. To this end, we propose a *parameter initializer* that manages ISP parameters of the first frame at the training time.

Parameter Initializer

As mentioned above, this module sets ISP parameters of the first frame at training time. The parameters are needed to be not always optimal values for the input to train how to improve. However, it still should be realistic values - learning parameter domain that will never be used at inference time is inefficient and makes the training difficult. Additionally, hand tuning the range for all parameters is difficult. Therefore we propose a method to choose realistic parameters automatically based on what kind of parameters was predicted as optimal for the second frames in the previous iterations by the controller as illustrated in Fig. 3. Our approach is to memorize the estimated parameters in the most recent M data in a buffer and randomly use the value at the first frame. We also consider sampling parameters from Gaussian distribution based on running mean and variance of the past parameters. This approach assume the parameters follow Gaussian distribution, but we are not sure it is true. Therefore, we expect our approach can generate more plausible initial parameters easily than other approaches which requires some assumptions or domain knowledge.

4. Evaluation

4.1. Datasets

Although DynamicISP can be applied to any recognition task, it is evaluated on detection tasks because of their wide usage. Two datasets are used to check the robustness in dif-




Figure 3. The proposed training method. The *parameter initializer* decides the ISP parameters for the first frame based on what kind of parameters were used for the second frames.

ferent settings. One is a human detection dataset [45]. This dataset aims to detect humans in any environment while training only with simple environmental images, as annotating difficult images is costly. Most of the test images are taken in challenging environments: dark, HDR, and heavy handshakes. They are taken with a RAW Bayer sensor. It has 18,880 easy images for the training set and 2,800 difficult images for the test set.

The other is LODDataset [14], which contains 2230 14-bit low-light RAW images with eight categories of objects. This dataset aims to detect multi-category objects in a low-light environment.

4.2. ISP functions

We implement five ISP functions, auto gain (AG), denoiser (DN), sharpener (SN), gamma tone mapping (GM), and contrast stretcher (CS), in a differentiable manner. The details are in the supplemental material. In the setting of single-function ISPs, GM is used as an ISP since it is known to improve image recognition accuracy the most [10]. In the setting of multi-function ISPs, the combination of functions is determined for each dataset after several experiments. It is known that the optimal combination of functions varies depending on the sensor, task, and environment [46]. In the future, this process could be automated by combining our method with a neural architecture search like ReconfigISP [46].

4.3. Ablation Studies on Human Detection

Various ablation studies are performed using the human detection RAW dataset [45]. We evaluate with the following setting if not specified. TTFNet [21] with ResNet18 [12] backbone is used as the detector. The entire pipeline is trained for 24 epochs with Adam optimizer [16] from randomly initialized weights using a cosine decay learning rate schedule (0.001 magnitudes) whose maximum and minimum learning rates are $1e-3$ and $1e-6$ with a linear warmup for the first 1,000 iterations [22]. The controller’s learning rate is multiplied by 0.1 from the base learning rate. After that, the model is additionally trained for 24 epochs with the same learning rate scheduling, but the ISP and controller

are frozen the same with NeuralAE [30]. This procedure slightly improves the accuracy from simple 48-epoch training. The input size is set to (576, 352, 3), and Rawgment [45] is applied to bridge the domain gap between the training data and challenging test data. During the training phase, the identical frame is input twice instead of using the continuous frame as described in Section 3.2 due to the rough frame rate of about 1 fps. In the evaluation phase, on the other hand, we try both sequential input and identical input. We set the channel size of the latent variable in the proposed controller as 256. The accuracy is evaluated with the average AP@0.5:0.95 [18] scores of four-time experiments with different random seeds.

Importance of Each Proposed Component

Table 1. Ablation studies about each proposed component. We add them one by one to the controller of NeuralAE [30]. The abbreviations of the proposed method are as follows; RO: *residual output format of parameters* whose default parameters are static, RO+: *residual output format of parameters* whose default parameters are learnable, PI: *parameter initializer*, and LU: *latent update style controller*. The metrics are AP@0.5:0.95 [%].

				ISP components	
PI	RO	RO+	LU	GM	DN+SN+GM+CS
				42.8	46.0
✓				45.8	-
✓	✓			48.6	-
✓	✓	✓		48.9	48.0
✓	✓	✓	✓	-	49.5

We add each proposed component one by one to the NeuralAE controller [30] to control two ISPs. One consisting of only GM and the other consisting of DN, SN, GM, and CS. The reason we do not use AG is that we tried to control AG, but for some reason, it failed on this dataset. The result is shown in Table 1.

Even when controlling only the GM, the complexity of the GM function makes it difficult to control with the NeuralAE method. The *parameter initializer* (PI) contributes significantly to alleviate the difficulty; without PI, the detector’s backbone has to deal with the pixel distributions before and after ISP, degrading detection performance. A detailed ablation study of the PI is shown in Table 2. We compare three types of PIs. One is a uniform sampling from $(p_{l,n,min}, p_{l,n,max})$; it learns how to control from all possible parameters. Another adds Gaussian noise to the running mean of the previously used parameters by tracking the running mean and running variance of used parameters in the second frame. The last approach is our best method, which memorizes the parameters for the 500 most recent data in a buffer and randomly adopts the value in the first frame.




Figure 4. The visualization result. We set an adequate confidence threshold per model as the precision@0.5 becomes 80% for a fair comparison. The darker bounding boxes are ground truth, and the brighter bounding boxes with captions are predictions.

It takes into account correlations between parameters and does not approximate a Gaussian distribution, thus covering more necessary and sufficient parameter regions. Our buffer memory automatically fits the search space to the minimum required, although it can still make the controller stable to the drastic environmental change. The buffer memory creates a situation where the environment is changed from day to night in a one-time step.

Table 2. Ablation studies about the *parameter initializer*.

	AP@0.5:0.95
None	42.8
uniform	44.2
running mean + Gaussian	44.7
buffer memory	45.8

The *residual output format of parameters* (RO) further eases the difficulty of controlling complex functions by making the controller output the difference from the averaging good parameters for the entire dataset. For the normal RO, the static ISP parameter obtained by differentiable tuning in Table 3 is used as a constant $p_{l,n}$. On the other hand, $p_{l,n}$ in RO+ is jointly optimized with the dynamic ISP control from scratch. The gain from RO to RO+ is thought to be because the averaging good parameters for the static ISP and DynamicISP is different. RO+ is practical as the parameters are tuned automatically.

For multi-layer control, the *latent update style controller* (LU) further improves the accuracy. All-at-one estimation of all the parameters in the multi-layers is difficult and the proposed LU successfully disentangles the problem setting.

Static Tuning V.S. Dynamic Control

The proposed dynamic control is compared to static tuning approaches. In this comparison, we use an ISP containing only a GM tone mapping function because tone mapping is known to have the biggest impact on machine vision [10]. One of the static tuning approaches for comparison is a grid search [45]. In this approach, the detector is trained per searched parameters. Since it takes a huge time to train the detector per parameter set of large search space, it searches for one parameter of the simplest gamma function $y = x^{\frac{1}{\gamma}}$. The other is a differentiable tuning approach to train the ISP and detector end-to-end [45]. Table 3 shows the effectiveness of our dynamic control. Although both static ISPs are optimized for machine vision, they are just optimal for the entire dataset and not optimal for each image.

Table 3. Static tuning V.S. dynamic control. The grid search method is the value described in Rawgnt performs[45] using a simple gamma. The diff. tuning denotes differentiable tuning of the static ISP parameters also described in Rawgnt [45].

	ISP components	
	GM	DN+SN+GM+CS
grid search [45]	45.3	-
diff. tuning [45]	48.3	45.1
dynamic control	48.9	49.5

Sequential Control Availability

The experiments above are conducted by inputting identical images twice instead of sequential input. We now check the

availability of sequential inference without twice input. In this setting, optimal parameters for the previous frame are used, resulting in possibly suboptimal. The result is shown in Table 4. Contrary to expectations, a slight improvement is observed, and it might come from the fact that the twice input case is evaluated under severe settings. The initial parameters are set as the moving average of the parameters in the training phase and are far from the optimal parameters per image. On the other hand, in the sequential settings, the previous parameters can be already decent values since consecutive frames are similar. From this experiment, efficient inference without input twice is proved to be possible. Additionally, since the dataset is captured at about 1 fps, and about half of it is taken with a shaken camera, consecutive frames may have some disparities. Inference at higher frame rates might further improve the performance of sequential inference.

Table 4. The availability of efficient sequential inference. The “input twice” is the same as all the experiments above. It uses the same frame for determining ISP’s parameters and detection. On the other hand, in “input sequentially,” optimal parameters for the previous frame is used for the next frame.

	ISP components	
	GM	DN+SN+GM+CS
input twice	48.9	49.5
input sequentially	49.1	49.6

Computational Efficiency

Table 5. The computational cost of the inference pipeline. C is the computational cost of the ISP. Any ISP configurations don’t change the computational cost of the controller very much; adding an ISP function increases the cost by only about $4e-5$ GFLOPS.

	computational cost [GFLOPS]			
	ISP	controller	detector	overall
w/o control	C	-	13.63	$13.63+C$
input twice	$2C$	0.02	15.99	$16.02+2C$
input sequentially	C	0.02	13.63	$13.65+C$

The computational efficiency is evaluated by profiling the FLOPS of each component. As listed in Table 5, the computational cost of our ISP controller is lightweight enough to be ignored compared with the detector despite our choice of a lightweight TTFNet detector with a ResNet18 backbone. Although our proposed LU has a complicated computational graph as shown in Fig. 2, the computational cost does not change very much (an increase of about $4e-5$ GFLOPS per each ISP function) thanks to the lightweight design with 1-D latent variable manipulations. The computational cost of the detector in the input twice

setting is not doubled from other settings because the whole part of the calculation for the first input is not needed. If ISP functions are lightweight classical functions, the dynamic ISP control can be efficient even as a single image detector with the inputting twice setting.

Comparison with State-of-the-arts

Various evaluations are performed against the state-of-the-art. The latency is measured on an RTX2080 GPU. The result is shown in Table 6.

One category for comparison is DNN-based encoder-decoder ISPs. The problem with the DNN ISP is that there are no input and ground truth pairs to train them. As pointed out in [33], creating clean ground truth for challenging noisy/blurry environment data is difficult. So, three kinds of DNN ISPs are tried. One is UNet(s); an UNet [36] based ISP, and the ISP and detector are trained end-to-end from scratch. It means that the ISP is trained only with the detection loss. The second is AW+sGM+UNet(f); a fixed state-of-the-art UNet ISP trained for dark environments [33] is used following auto-white balance (AWB) and simple gamma (sGM). The AWB and sGM are identical to the preprocessing pipeline of the trained UNet. In this setting, the UNet ISP is trained with paired data for dark environments proposed in [33]. The third is AW+sGM+UNet(p); end-to-end training is performed from the pre-trained UNet ISP’s weight to tune it to the target sensor and downstream detector.

The other category is classical function-based ISPs. As for NeuralAE, we do not convert the images to those taken with lower-quality sensors as the original paper did to make a fair comparison and use our ISP functions because the details are missing. Comparing DNN ISP and classical ISPs, the tuned classical ISPs outperform the existing state-of-the-art DNN ISPs in challenging environments despite the low computational cost. One of the reasons should be the lack of paired data for the target dataset to train the DNN ISPs, but it is difficult to create ground truth images for challenging RAW data. ISPs that do not require paired data and can be optimized with downstream task loss must be useful. As to NeuralAE [30], we try two of its proposed encoders, *Semantic Feature Branch* and *Global Feature Branch*. Our *Global Feature Branch* implementation uses the `histc()` function implemented in Pytorch [31] to generate multiple histograms of the image, but it is time-consuming. So we only use *Semantic Feature Branch* in our method. We also tried further tuning from the well-trained diff. tuning (GM) [30] weight using Hardware-in-the-loop [45], but it did not improve accuracy due to overfitting the training data. Although we use vanilla CMA-ES in the Hardware-in-the-loop due to the missing information about the modifications, it should be no problem since only

Table 6. Comparison with state-of-the-arts. C and H are the computational cost of the ISP and generating histograms of an image.

ISP		GFLOPS	#params [M]	latency [ms]	w/o detector [ms]	AP [%]
w/o ISP		13.63	14.20	7.0	0.0	32.8
DNN based	UNet(s)	21.96	21.96	13.8	6.8	28.7
	AWB+sGM+UNet(f) [33]	21.96+C	21.96	14.2	7.2	17.3
	AWB+sGM+UNet(p) [33]	21.96+C	21.96	14.2	7.2	42.0
classical function based	sRGB (black box ISP)	13.63	14.20	-	-	38.4
	diff. tuning (ISP=DN+SN+GM+CS) [45]	13.63+C	14.20	25.8♠	18.8♠	45.1
	NeuralAE (ISP=DN+SN+GM+CS) [30]	13.65+C	15.66	26.0♠	19.0♠	45.2
	NeuralAE (ISP=GM)♣ [30]	13.65+C+H	15.66	8.4	1.4	47.0
	Hardware-in-the-loop [27]	13.63+C	14.20	8.1	1.1	47.1
	diff. tuning (ISP=GM) [45]	13.63+C	14.20	8.1	1.1	48.3
	NeuralAE (ISP=GM)† [30]	13.65+C+H	15.66	18.3	11.3	48.5
dynamic	NeuralAE (ISP=GM) [30]	13.65+C	15.66	8.4	1.4	48.6
	ours (ISP=GM)	13.65+C	15.66	8.3	1.3	49.1
	ours (ISP=DN+SN+GM+CS)	13.65+C	15.76	27.7♠	20.7♠	49.6

♣: The original augmentation in stead of Rawgment[45] is used.

†: Full version of NeuralAE including *Global Feature Branch*, which is a histogram-based encoder to support *Semantic Feature Branch*.

♠: Our implementations of DN and SN have a lot of duplicate calculations and are not optimized.

three parameters in the GM function are the target for the optimization.

Comparing ours with both categories, our method achieves the best accuracy in spite of the low computational cost. Especially, compared with diff. tuning, ours further improved detection accuracy with a little computational gain. While our computational cost and inference speed are almost unchanged compared to static ISP (diff. tuning), the number of parameters is increased. This is because the fully connected layers whose channel sizes are 1024 in the *Semantic Feature Branch* have many parameters despite a few computational costs. If the target device is constrained by the number of parameters, we can reduce their channel sizes. It is not a bottleneck, but an enough large value of 256 is also set for the channel size for latent variable in LU since the computational cost hardly increases. Although NeuralAE also controls the exposure function, it is a simple function. Instead, ours enables the control of complicated functions and yields more suitable images for the downstream detector like Fig. 4.

4.4. Low Light Image Recognition

We evaluate on the public LODDataset [14]. The experimental settings are almost identical to [14] using CenterNet [50] as a detector. The only one difference is that we apply wider contrast augmentation from 1/100 to 1/10 since the dataset is captured with 1/100 to 1/10 times of default ISO and our input are not normalized with the ISO factor. Following [14], we use simulated RAW-like images converted from COCO Dataset [18] using UPI [2] as a training dataset and use real RAW images as a test dataset. AG, GM, and CS are used for the ISP as this combination is deemed effective

for LODDataset. It might be because this dataset is not so much darker than the darkest image in the human detection dataset.

As shown in Table 7, our method achieves significant improvement from the state-of-the-arts. The dynamic control may mitigate the domain gap between real RAW images and simulated RAW-like images. Because the control is based on what the detector felt and reacted to the input image, it automatically self-adjusts and closes the domain gap. Our method outperformed NeuralAE by a large margin by enabling the control of complex ISP functions. More results are listed in the supplemental material.

Table 7. Evaluation on LODDataset trained with simulated RAW-like data converted from COCO Dataset [18]. Note that NeuralAE and ours are compared with the same ISP combinations. NeuralAE always applies AE outside ISPs, but it works samely with AG on the software computation.

	AP@0.5:0.95
as is	14.3
SID [4]	16.2
Zero DCE [9]	21.6
REDI [17]	25.4
H. Yang et. al. [14]	28.8
diff. tuning (ISP=GM) [45]	33.0
NeuralAE (ISP=GM)[30]	33.7
NeuralAE (ISP=GM+CS)[30]	30.6
ours (ISP=GM)	43.6
ours (ISP=AG+GM+CS)	44.0

5. Conclusion

We propose DynamicISP, a method to dynamically control parameters of the ISP functions per image, which is beneficial to recognition in various environments. Although they are classical functions, the dynamic control compensates expressive power. Our controller is based on how the downstream recognition model reacted and desired at the previous time step. The control is realized with the proposed *residual output format of parameters*, *latent update style controller*, and *parameter initializer*. Finally, we achieve state-of-the-art detection results with efficient computing costs for detection in various environments. We believe DynamicISP will enhance many other computer vision tasks, such as image classification, keypoint detection, and segmentation. Furthermore, the proposed controller might enhance other modal tasks like sound recognition and natural language processing.

References

- [1] Mahmoud Afifi, Jonathan T Barron, Chloe LeGendre, Yun-Ta Tsai, and Francois Bleibel. Cross-camera convolutional color constancy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1981–1990, 2021. [2.1](#)
- [2] 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*, pages 11036–11045, 2019. [4.4](#)
- [3] Mark Buckler, Suren Jayasuriya, and Adrian Sampson. Reconfiguring the imaging pipeline for computer vision. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 975–984, 2017. [1](#)
- [4] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3291–3300, 2018. [7](#)
- [5] Liangyu Chen, Xin Lu, Jie Zhang, Xiaojie Chu, and Chengpeng Chen. Hinet: Half instance normalization network for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 182–192, 2021. [2.1](#)
- [6] Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, and Zicheng Liu. Dynamic convolution: Attention over convolution kernels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11030–11039, 2020. [1](#), [2.3](#), [3.1](#)
- [7] Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, and Zicheng Liu. Dynamic relu. In *European Conference on Computer Vision*, pages 351–367. Springer, 2020. [2.3](#)
- [8] Steven Diamond, Vincent Sitzmann, Frank Julca-Aguilar, Stephen Boyd, Gordon Wetzstein, and Felix Heide. Dirty pixels: Towards end-to-end image processing and perception. *ACM Transactions on Graphics (TOG)*, 40(3):1–15, 2021. [1](#), [2.1](#), [2.2](#)
- [9] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1780–1789, 2020. [2.2](#), [7](#)
- [10] Patrick Hansen, Alexey Vilkin, Yury Krustalev, James Imber, Dumidu Talagala, David Hanwell, Matthew Mattina, and Paul N Whatmough. Isp4ml: The role of image signal processing in efficient deep learning vision systems. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 2438–2445. IEEE, 2021. [1](#), [4.2](#), [4.3](#)
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015. [2.3](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [4.3](#)
- [13] Luis V Hevia, Miguel A Patricio, José M Molina, and Antonio Berlanga. Optimization of the isp parameters of a camera through differential evolution. *IEEE Access*, 8:143479–143493, 2020. [2.1](#)
- [14] Yang Hong, Kaixuan Wei, Linwei Chen, and Ying Fu. Crafting object detection in very low light. In *Proceedings of the British Machine Vision Virtual Conference*, 2021. [2.2](#), [4.1](#), [4.4](#), [7](#)
- [15] 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. [2.2](#)
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [4.3](#)
- [17] Mohit Lamba and Kaushik Mitra. Restoring extremely dark images in real time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3487–3497, 2021. [7](#)
- [18] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. [4.3](#), [4.4](#), [7](#)
- [19] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018. [2.1](#)
- [20] Shuai Liu, Chaoyu Feng, Xiaotao Wang, Hao Wang, Ran Zhu, Yongqiang Li, and Lei Lei. Deep-flexisp: A three-stage framework for night photography rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1211–1220, 2022. [1](#), [2.1](#), [2.2](#)
- [21] Zili Liu, Tu Zheng, Guodong Xu, Zheng Yang, Haifeng Liu, and Deng Cai. Training-time-friendly network for real-time

- object detection. In *proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 11685–11692, 2020. 4.3
- [22] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 4.3
- [23] Feifan Lv, Feng Lu, Jianhua Wu, and Chongsoon Lim. Mblen: Low-light image/video enhancement using cnns. In *BMVC*, volume 220, page 4, 2018. 2.2
- [24] Long Ma, Tengyu Ma, Risheng Liu, Xin Fan, and Zhongxuan Luo. Toward fast, flexible, and robust low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5637–5646, 2022. 2.2
- [25] Kristina Monakhova, Stephan R Richter, Laura Waller, and Vladlen Koltun. Dancing under the stars: video denoising in starlight. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16241–16251, 2022. 2.1
- [26] Igor Morawski, Yu-An Chen, Yu-Sheng Lin, Shusil Dangi, Kai He, and Winston H Hsu. Genisp: Neural isp for low-light machine cognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 630–639, 2022. 1, 2.1, 2.2
- [27] Ali Mosleh, Avinash Sharma, Emmanuel Onzon, Fahim Mannan, Nicolas Robidoux, and Felix Heide. 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*, pages 7529–7538, 2020. 1, 2.1, 6
- [28] Jun Nishimura, Timo Gerasimow, Rao Sushma, Aleksandar Sutic, Chyuan-Tyng Wu, and Gilad Michael. Automatic isp image quality tuning using nonlinear optimization. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 2471–2475. IEEE, 2018. 2.1
- [29] Taishi Ono, Yuhi Kondo, Legong Sun, Teppei Kurita, and Yusuke Moriuchi. Degree-of-linear-polarization-based color constancy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19740–19749, 2022. 2.1
- [30] Emmanuel Onzon, Fahim Mannan, and Felix Heide. Neural auto-exposure for high-dynamic range object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7710–7720, 2021. 1, 2.1, 2.3, 3.1, 3.1, 3.2, 4.3, 1, 4.3, 4.3, 4.3, 6, 7
- [31] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 4.3
- [32] G Pavithra and Bhat Radhesh. Automatic image quality tuning framework for optimization of isp parameters based on multi-stage optimization approach. *Electronic Imaging*, 2021(9):197–1, 2021. 2.1
- [33] Abhijith Punnappurath, Abdullah Abuolaim, Abdelrahman Abdelhamed, Alex Levinstein, and Michael S Brown. Day-to-night image synthesis for training nighttime neural isps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10769–10778, 2022. 1, 2.1, 2.2, 4.3, 4.3, 6
- [34] Erik Reinhard, Wolfgang Heidrich, Paul Debevec, Sumanta Pattanaik, Greg Ward, and Karol Myszkowski. *High dynamic range imaging: acquisition, display, and image-based lighting*. Morgan Kaufmann, 2010. 1
- [35] Nicolas Robidoux, Luis E Garcia Capel, Dong-eun Seo, Avinash Sharma, Federico Ariza, and Felix Heide. End-to-end high dynamic range camera pipeline optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6297–6307, 2021. 1, 2.1
- [36] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. 4.3
- [37] Eli Schwartz, Alex Bronstein, and Raja Giryes. Isp distillation. *arXiv preprint arXiv:2101.10203*, 2021. 2.2
- [38] Stanley S Stevens. On the psychophysical law. *Psychological review*, 64(3):153–181, 1957. 1
- [39] Ethan Tseng, Felix Yu, Yuting Yang, Fahim Mannan, Karl ST Arnaud, Derek Nowrouzezahrai, Jean-François Lalonde, and Felix Heide. Hyperparameter optimization in black-box image processing using differentiable proxies. *ACM Trans. Graph.*, 38(4):27–1, 2019. 1, 2.1
- [40] Zhengzhong Tu, Hossein Talebi, Han Zhang, Feng Yang, Peyman Milanfar, Alan Bovik, and Yinxiao Li. Maxim: Multi-axis mlp for image processing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5769–5780, 2022. 2.1
- [41] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. *arXiv preprint arXiv:1808.04560*, 2018. 2.2
- [42] Jay Whang, Mauricio Delbracio, Hossein Talebi, Chitwan Saharia, Alexandros G Dimakis, and Peyman Milanfar. Deblurring via stochastic refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16293–16303, 2022. 2.1
- [43] Chyuan-Tyng Wu, Leo F Isikdogan, Sushma Rao, Bhavin Nayak, Timo Gerasimow, Aleksandar Sutic, Liron Ainkedem, and Gilad Michael. Visionisp: Repurposing the image signal processor for computer vision applications. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 4624–4628. IEEE, 2019. 2.1
- [44] Brandon Yang, Gabriel Bender, Quoc V Le, and Jiquan Ngiam. Condconv: Conditionally parameterized convolutions for efficient inference. *Advances in Neural Information Processing Systems*, 32, 2019. 1, 2.3, 3.1
- [45] Masakazu Yoshimura, Junji Otsuka, Atsushi Irie, and Takeshi Ohashi. Rawgmont: Noise-accounted raw augmentation enables recognition in a wide variety of environments. *arXiv preprint arXiv:2210.16046*, 2022. 2.1, 4.1, 4.3, 4.3, 4.3, 3, 4.3, 6, 7

- [46] Ke Yu, Zexian Li, Yue Peng, Chen Change Loy, and Jinwei Gu. Reconfigisp: Reconfigurable camera image processing pipeline. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4248–4257, 2021. [1](#), [2.1](#), [4.2](#)
- [47] Meina Zhang, Yingying Fang, Guoxi Ni, and Tiejong Zeng. Pixel screening based intermediate correction for blind deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5892–5900, 2022. [2.1](#)
- [48] Yi Zhang, Dasong Li, Ka Lung Law, Xiaogang Wang, Hongwei Qin, and Hongsheng Li. Idr: Self-supervised image denoising via iterative data refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2098–2107, 2022. [2.1](#)
- [49] 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. [2.2](#)
- [50] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. [4.4](#)