

# LW-ISP: A Lightweight Model with ISP and Deep Learning

Hongyang Chen<sup>1</sup>  
chenhy@stu.xjtu.edu.cn

Kaisheng Ma<sup>2</sup>  
kaisheng@mail.tsinghua.edu.cn

<sup>1</sup> Xi'an Jiaotong University  
Xi'an, China

<sup>2</sup> Tsinghua University  
Beijing, China

---

## Abstract

The deep learning (DL)-based methods of low-level tasks have many advantages over the traditional camera in terms of hardware prospects, error accumulation and imaging effects. Recently, the application of deep learning to replace the image signal processing (ISP) pipeline has appeared one after another; however, there is still a long way to go towards real landing. In this paper, we show the possibility of learning-based method to achieve real-time high-performance processing in the ISP pipeline. We propose LW-ISP, a novel architecture designed to implicitly learn the image mapping from RAW data to RGB image. Based on U-Net architecture, we propose the fine-grained attention module and a plug-and-play upsampling block suitable for low-level tasks. In particular, we design a heterogeneous distillation algorithm to distill the implicit features and reconstruction information of the clean image, so as to guide the learning of the student model. Our experiments demonstrate that LW-ISP has achieved a 0.38 dB improvement in PSNR compared to the previous best method, while the model parameters and calculation have been reduced by  $23\times$  and  $81\times$ . The inference efficiency has been accelerated by at least  $15\times$ . Without bells and whistles, LW-ISP has achieved quite competitive results in ISP subtasks including image denoising and enhancement.

## 1 Introduction

In recent years, smartphones have increasingly dominated daily photos. With the emergence of advanced applications such as autonomous driving [6, 5], high-speed continuous shooting [6] and 4K recording [7], the importance and requirements for cameras are increasing gradually. The image signal processing (ISP) [2, 4, 4, 5] is used to receive and process the raw signal of the sensor during the entire process of camera imaging, which has a decisive effect on the quality of the image. As mobile devices begin to adapt to powerful hardware with an ISP system [4], the resolution has been greatly increased. However, small sensors and relatively compact lenses have led to the loss of detail and high noise levels, and the current ISP system still fails to solve these problems completely. Moreover, as Deep Neural Networks (DNNs) achieve performance that surpasses conventional algorithms in the tasks of image classification [8, 5], speech recognition [7, 2] and other fields [4, 9, 9], the combination of DNNs and ISP has been brought to the fore.

Traditional ISPs and DL-based solutions still face various challenges. As the special hardware in the camera dedicated to image processing tasks, ISP can solve many low-level and global image processing tasks in proper order, such as demosaicing, white balance, exposure correction, and gamma correction. In the design of the traditional ISP system, aforementioned tasks are well researched independently, without considering its subsequent impact, which may lead to the accumulation of errors in the entire processing pipeline. That means that the overall process will be affected by error propagation from stage to stage. For instance, early demosaicing artifacts may be amplified by image sharpening or misalignment of different exposures [23]. At present, learning ISP pipeline promotes a novel direction of research aiming at replacing the current tedious and expensive handcrafted ISP solutions with data-driven learned ones capable of surpassing them in terms of image quality. The advantages of learning-based methods are that they can implicitly learn the statistical information of natural images and allow joint solutions for multiple tasks. However, the limited research [26, 30, 45, 47] mainly focuses on the improvement of objective indicators or only targets closely related tasks. These methods customarily require a higher computational overhead, which is challenging to be taken into account by the application.

In this paper, we propose LW-ISP to replace the entire ISP pipeline, achieving an effective balance of processing efficiency and performance. The first step is to design a tiny U-Net as the base model and the Fine-grained Attention Module (FGAM) to reconstruct the overall information during down-sampling processes. The second step towards incorporating contextual information into the upsampling blocks to preserve realistic details from RAW inputs, we design a plug-and-play Contextual Complement Upsampling Block (CCB). Finally, we design the heterogeneous distillation algorithm to train the teacher model based on the target data and then distill the clean features from the teacher model to the student.

Our experiments demonstrate the superiority of the LW-ISP on the ISP dataset and its subtasks. In a large-scale learning setup, our approach achieves a performance exceeding SOTA methods on the ISP’s largest dataset and dramatically speeds up the inference process (model parameters are reduced by 23 times and calculations are reduced by 81 times), as shown in Fig. 1. Furthermore, our training techniques improve transfer performance on a suite of ISP downstream subtasks such as image denoising and image enhancement. We recommend that practitioners use this simple architecture as a baseline for future research.

To sum up, the contribution of this paper can be summarized as follows.

- We design a novel lightweight LW-ISP model to replace the entire ISP pipeline. We evaluate the performance of LW-ISP on the ISP pipeline and different subtasks. Plainly speaking, LW-ISP not only outputs processed images with higher subjective and objective quality (PSNR: 0.38dB), but also takes less inference time (15×).
- We propose a fine-grained attention module to reconstruct the overall information and determine a more reasonable upsampling block in low-level image processing.

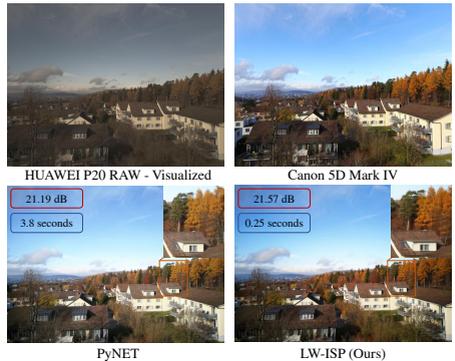


Figure 1: Example set of images from ISP dataset (Zurich [30]). The number in the red box represents the PSNR value, and the blue box indicates the test time under a single NVIDIA Tesla V100.

- We propose a heterogeneous distillation training algorithm to distill the spatial structure information and global information of the environment from the teacher model to the student model.

## 2 Related Work

**DL-based ISP Subtasks.** Deep learning (DL)-based methods have achieved considerable success in image preprocessing subtasks, including demosaicing [36], denoising [12], deblurring [63] and super-resolution [68], which all have achieved performance beyond conventional algorithms. Studies have shown that even when operating outside of a supervised learning mechanism, DNNs are proficient in generating high-quality images [60]. Recently, researchers denoise on RAW images in order to avoid the effect of ISP [65]. Brooks *et al.* creatively inversely transformed the color image and used it for the training [9]. Contrary to the conventional methods of independently solving ISP subtasks, DL-based methods allow multiple tasks to be jointly solved, which has great potential to reduce the computational burden. However, existing schemes require more calculations.

**DL-based ISP Pipeline.** The application of DNNs to solve the ISP pipeline has gradually attracted attention. As the first attempt of ISP with DL, DeepISP [47] divided the framework into high-dimensional and low-dimensional feature extraction parts, which perform local and global learning, respectively. Nevertheless, this method only considers two tasks of image demosaicing and denoising. Ratnasingam *et al.* reconstructed RGB images into RAW images to obtain a large number of training images [45]. PyNET [30] focused on the mobile camera ISP pipeline and processed images from five different levels to obtain higher quality information. However, the processing time in the CPU mode is as high as 100 seconds. Moreover, some previous works [9, 13, 15, 26, 33, 35, 43] in AIM 2020 Challenge [29, 65] and Mobile AI 2021 Challenge [31] have achieved appealing results. CSANet [26] designed a double attention structure for mobile ISP while showing significant differences on other datasets. Moreover, the replacement of ISP is also based on HDR [9] and extreme low-light [61]. CameraNet [35] defined ISP as restoration and enhancement subtasks, which are learned in two stages. More

In terms of replacing an existing handcrafted ISP pipeline, these methods pay too much attention to the improvement of objective indicators or can only accomplish closely related tasks. In this paper, we abandon the redundant architecture to replace the model of the entire ISP pipeline. Moreover, we try to surpass the work of predecessors in ensuring a lightweight model and realizing the real coordinated development of ISP and AI vision.

## 3 Proposing LW-ISP

In this section, we present an overview of our proposed LW-ISP, as illustrated in Fig. 2. Instead of naively adopting multi-scale or serial modular architecture to process RAW input, we take advantage of sophisticated structure and training strategies to achieve lightweight and high performance. We design a tiny U-Net as the backbone, which only contains 24 layers of convolution. The bottom half of Fig. 2 shows our proposed end-to-end image preprocessing network, LW-ISP, which is composed of two stages. The first stage progressively downsamples feature maps at different levels to accelerate the computation. The second stage further concatenates the processed global vector with the tensor of the same

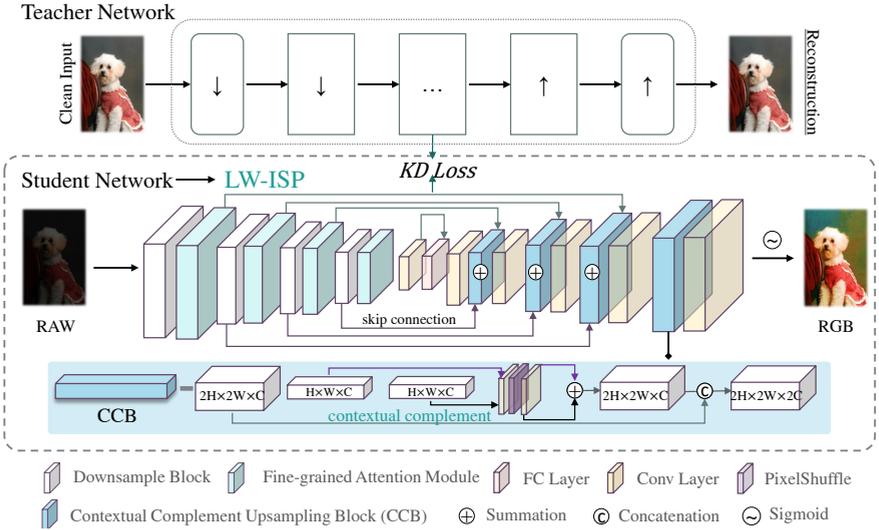


Figure 2: The overview of our method LW-ISP. The bottom half is the main architecture, which receives RAW input and execution feature reconstruction.

size in the first half of the network through the symmetric skip connection. *More details and codes about our architecture can be found in the supplementary material.*

### 3.1 Fine-Grained Attention Module

As LW-ISP performs feature extraction on the RAW input in the first stage, more effective and discernible features need to be passed on. Existing attention fusion structures [27, 56] can reconstruct pixels in high-dimensional spaces. Furthermore, some research has been proposed to implement the attention mechanism in low-level vision but exhibited completely different negative effects in the ISP task [24, 50].

To address this issue, we propose the *Fine-grained Attention Module (FGAM)* suitable for ISP, as shown in Fig. 3. The feature recalibration is achieved by performing the attention of the channel and spatial dimensions in parallel and then indirectly fusing the intermediate features. As to channel attention, given a feature map  $H \times W \times C$ , the squeeze operation applies global average pooling (GAP) and two convolution layers followed by sigmoid gating to generate activation vector  $1 \times 1 \times C$ . As for spatial attention, GAP and max pooling (GMP) operate on feature map  $H \times W \times C$  along the channel dimension and concatenate the outputs to yield a feature map  $H \times W \times 2$ . Then passing through a convolution layer and sigmoid to obtain the attention map  $H \times W \times 1$ . Specifically, in order to prevent the attention mechanism from suppressing the characteristics of irrelevant regions, we find that the direct fusion of spatial and channel attention should be avoided. In the FGAM, these two attention maps separately add to the input feature and then concatenate across channel dimension.

### 3.2 Contextual Complement Upsampling Block

The *Contextual Complement Upsampling Block (CCB)* is composed of a contextual separation module for adaptive high frequency decomposition in the feature space, followed by the CCB-Core that fuses the corresponding size ( $H \times W$ ) features of the previous stage. Unlike PyNET [30] that hand over the upsampling operation to the conventional method (bilinear

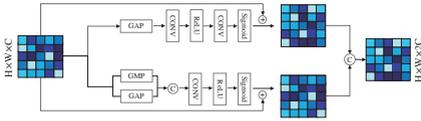


Figure 3: The fine-grained attention module (FGAM) combines channel attention (upper) and spatial attention (lower) with the input feature and then concatenate together.

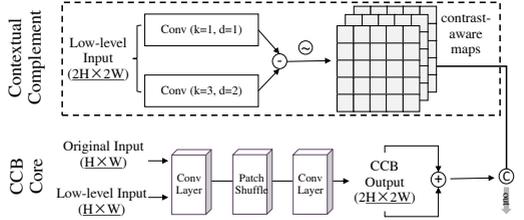


Figure 4: The contextual complement upsampling block (CCB).

interpolation) to solve, we switch to PixelShuffle [18] and design CCB-Core to suppress the loss of image information during zooming.

**CCB-Core.** As shown by the dashed box in Fig. 2 and Fig. 4, in addition to the original features to be sampled, the input of CCB-Core also has the corresponding features ( $2H \times 2W$  for contextual decomposition and  $H \times W$  for feature complementation) of the previous stage. The CCB-Core first performs a sub-pixel convolution on input features ( $H \times W$  features from the first and second stages) to obtain global- and local-features. It should be noted that the convolution operations before and after PixelShuffle are used for channel dimension adjustment and fine-tuning, respectively. Subsequently, CCB-Core fuses the obtained features by residual learning to derive coarse high-resolution features. The operation of CCB-Core can be formulated as follows:

$$\mathcal{O}_{FAU} = \mathcal{P}(\mathcal{O}_{FBU}^1) + \mathcal{P}(\mathcal{O}_{FBU}^2), \quad (1)$$

where  $\mathcal{P}$  and  $+$  stand for the functions of the sub-pixel convolution and residual learning, respectively.  $FAU/FBU$  are short for the feature after/before upsampling, and the numbers represent the stage of the feature.

**Contextual Complement.** This part is to learn contrast-aware features for image decomposition. To select adaptive contextual information, we first use two groups of dilated convolutions (with kernel size&dilation rate of 1&1 and 3&2), denoted as  $f_{d1}$  and  $f_{d2}$ , to extract features in different receptive fields. The effectiveness of this process will be verified in Section 4.3. We then compute a contrast-aware map between the two feature maps as:

$$C_l = \text{sigmoid}(f_{d1}(x_{in}) - f_{d2}(x_{in})), \quad (2)$$

where  $C_l$  indicates the pixel-wise relative contrast information.  $C_l$  will eventually be concatenated to the  $\mathcal{O}_{FAU}$  of CCB-Core to complete the contextual complement.

### 3.3 Heterogeneous Distillation Algorithm

In knowledge distillation, although the source data distribution and processing dimensions of several networks may be different, target models can still be imitated by knowledge distillation on the target data [24, 25]. For ISP, its reconstruction process requires more hidden features and spatial structure information. In this paper, we propose a heterogeneous distillation algorithm, as shown in Fig. 2. The teacher will learn a wealth of intrinsic attributes from the clean inputs to assist the student network.

For the teacher network, we continue to adopt the basic structure of LW-ISP. The only difference lies in the lack of an upsample block, due to the fact that RAW inputs will be half

of the normal output after pre-processing. The teacher model will learn a clean mapping from the ground truth of the training data. The student model, LW-ISP, will be supervised during the training process, and no burden will be added during inference. We supervise the intermediate features and compute the feature similarity to control the imitation learning from the teacher to student. The determination of the position of the intermediate feature requires to be carefully selected. Besides, the heterogeneous distillation can also be used to extend the ISP pipeline to more fine-grained tasks. Experiments demonstrate that when applying the heterogeneous distillation, the position of the intermediate feature needs to be carefully selected, and the supervision provided in the first stage even harms the performance. We cherry-pick the output of the last three upsample blocks.

### 3.4 Loss Function for LW-ISP

Based on the above distillation algorithm and full ISP’s requirements for local and global correction and perceptual quality, we design multiple loss functions for training. The overall loss of LW-ISP can be formulated as:

$$\mathcal{L}_{\text{overall}} = \mathcal{L}_r + \alpha \cdot \mathcal{L}_s + \beta \cdot \mathcal{L}_d, \quad (3)$$

where  $\alpha$  and  $\beta$  are two hyper-parameters to balance the magnitude of  $\mathcal{L}_s$  and  $\mathcal{L}_d$ . The sensitivity study and ablation study have been shown in Section 4.3.

**Reconstruction Loss.** Given a training sample  $I$ , the model result and ground truth can be denoted as  $f(I)$  and  $J$ . To obtain the reconstruction result, we adopt the mean absolute error (MAE) to measure the difference:

$$\mathcal{L}_r = |f(I) - J|. \quad (4)$$

**Structural Loss.** The multi-scale structural similarity (MS-SSIM [53]) loss is used here to increase the dynamic range of the reconstructed photos:

$$\mathcal{L}_s = 1 - MS \cdot SSIM(f(I), J). \quad (5)$$

**Distillation Loss.** The teacher network transmits the implicit information of the clean image to the student through the intermediate representations. Denote  $F_m$  to be the feature maps of the  $m_{th}$  layer of the model, the same is true for  $n$ . The distillation loss (KD loss) is formulated as:

$$\mathcal{L}_d = \sum_{(m,n) \in C} \mathcal{L}_2(F_m^s(I), (F_n^t(J))), \quad (6)$$

where  $\mathcal{L}_2$  is the  $L_2$ -norm loss and  $C$  is a set of candidate pairs of feature locations. In this work,  $m$  and  $n$  are set to be consistent. The superscript  $t$  and  $s$  denote the teacher model and student model, respectively.

### 3.5 Loss for Teacher Network

To learn an effective feature representation from the teacher model, we design the following loss function:

$$\mathcal{L}_T = (g(J) - J)^2 + \gamma \cdot \mathcal{L}_s, \quad (7)$$

where  $g$  is the transform function and  $J$  is the clean image.  $\gamma$  is a hyper-parameter to balance the magnitude.

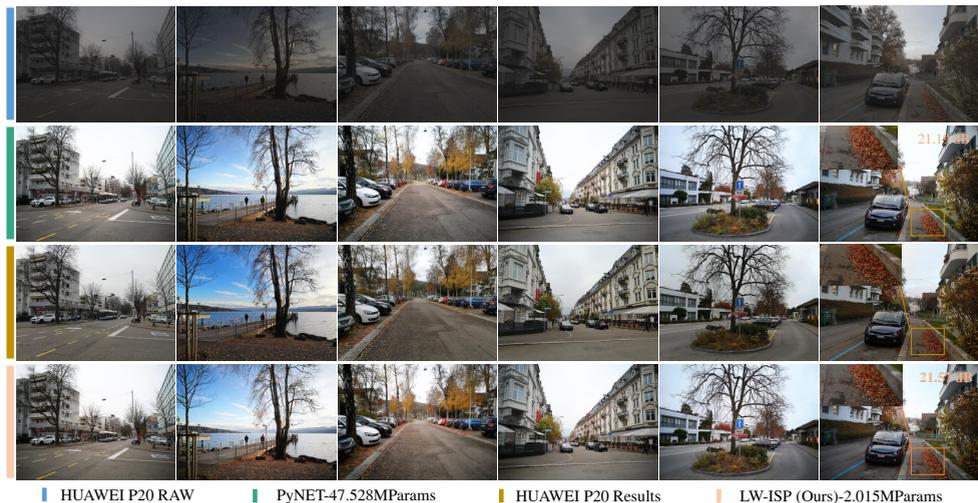


Figure 5: Sample visual results obtained with the proposed LW-ISP architecture (best zoomed on screen). These four lines represent the output of HUAWEI P20 sensor, state-of-the-art model PyNET, HUAWEI P20 camera and our method, respectively. Besides, the two numbers represent the evaluation results of the model with the PSNR.

Method	SRCNN[ <a href="#">30</a> ]	SRGAN[ <a href="#">30</a> ]	DPED[ <a href="#">30</a> ]	U-Net[ <a href="#">30</a> ]	Pix2Pix[ <a href="#">30</a> ]	SPADE[ <a href="#">30</a> ]	NAFNet[ <a href="#">30</a> ]	PyNET[ <a href="#">30</a> ]	LW-ISP
PSNR ( $\uparrow$ )	18.56	20.06	20.67	20.81	20.93	20.96	21.12	<u>21.19</u>	<b>21.57</b>
MS-SSIM ( $\uparrow$ )	0.8268	0.8501	0.8560	0.8545	0.8532	0.8586	0.8613	<u>0.8620</u>	<b>0.8622</b>
LPIPS ( $\downarrow$ )	0.385	0.257	0.343	0.257	0.208	0.209	0.194	<u>0.194</u>	<b>0.160</b>

Table 1: Comparison experiment results (PSNR/MS-SSIM/LPIPS) on Zurich Dataset [[30](#)] (numbers in bold are the best).  $\uparrow$  denotes that the upward trend corresponds to better performance, and  $\downarrow$  denotes the downward trend.

## 4 Experiment

When learning RAW-to-RGB mapping with deep learning methods, we refer to it as **smart ISP**. If the smart ISP is to be applied towards landing application, the first thing to solve is the need for real-time inference and high imaging performance. In this section, we evaluate the effectiveness of our method on Zurich RAW to RGB [[30](#)] (Zurich for short) dataset, which is currently the largest ISP dataset. The effects of image denoising and enhancement are also evaluated on the SIDD [[11](#)], DND [[12](#)] and LoL [[14](#)] datasets.

### 4.1 Results on Zurich Dataset

**Effectiveness of LW-ISP.** Due to limited deep learning research on the ISP pipeline, we compare different image preprocessing architectures, including the existing state-of-the-art method PyNET [[30](#)]. We adopt the mainstream image quality evaluation indicators PSNR, MS-SSIM and LPIPS. Table 1 shows the quantitative performance of the proposed method on the real RAW to RGB mapping problem. It is obvious that LW-ISP outperforms other state-of-the-arts with the gain of at least 0.38dB in terms of PSNR and performs well on MS-SSIM. Fig. 1 and Fig. 5 compare the visual effects of LW-ISP with the previous best model PyNET and HUAWEI P20 after processing different RAW images. It is observed

Model	Lightweight [50]	HERN [50]	CameraNet [50]	AWNet [50]	Pynet-ca [50]	LW-ISP (Ours)
PSNR (dB)	21.28	21.30	21.35	21.40	21.50	<b>21.57</b>
Params.(M)	31.56	39.64	26.53	55.70	56.89	<b>2.01</b>

Table 2: Comparison with SOTA methods (PSNR and number of parameters) on Zurich Dataset [50] (numbers in bold are the best).

Model	Number of Parameters	FLOPS			PSNR (dB)
		(224,224)	(960,960)	(1440,1984)	
SPADE [44]	97, 480, 899	191.31G	3.16T	10.89T	20.96
PyNET [50]	47, 554, 738	342.698G	5.72T	19.513T	21.19
LW-ISP w/o FGAM	<b>1, 660, 777</b>	<b>3.441G</b>	<b>63.198G</b>	<b>195.914G</b>	<u>21.40</u>
LW-ISP w/ FGAM	2, 014, 681	4.234G	69.198G	211.914G	<b>21.57</b>

Table 3: Comparison of the number of the parameters and FLOPs (floating point operations) between our proposed method and the state-of-the-art methods. Note that LW-ISP w/o FGAM does not adopt heterogeneous distillation and SPADE [44] is a linear architecture with basically no long-distance cross-layer connection.

that the images taken by HUAWEI P20 are generally dark and over-render the sky and other backgrounds. In contrast, the results of LW-ISP are more in line with the real characteristics and can generate better details, which are verified by objective metrics. The recently emerging NAFNet [44] attempts to design a simple baseline for the field of image restoration, even removing nonlinear activation functions. Our method surpasses the state-of-the-art backbone (NAFNet) in low-level vision. We believe that the dark characteristics of the RAW data itself require more refined processing.

**Efficiency of LW-ISP.** *How much loss in quality is tolerable for the increase in speed?* Our model aims to achieve an effective balance between processing efficiency and algorithm performance to promote the development of the smart ISP. Roughly, we provide LW-ISP w/o FGAM and LW-ISP w/ FGAM in order to provide more options. LW-ISP w/o FGAM means that the attention module FGAM is not added to LW-ISP. As shown in Table 3, compared to the SOTA PyNET [50], LW-ISP can achieve up to 28.6 times of compression, and adding FGAM can still reduce it by 23 times. The results also present the floating point operations (FLOPs) under different resolution conditions. The calculation of the latest LW-ISP is 81 times less than PyNET. What is surprising is that LW-ISP can still achieve better performance than the previous best method with a minimum acceleration of 15 times. When testing on NVIDIA Tesla V100 GPU, LW-ISP takes 0.25 seconds to process 12MP photo ( $2944 \times 3958$  pixels), while PyNET takes 3.8 seconds. *The memory usage and other params of our model are presented in the supplementary.*

**Comparison with SOTA Methods.** To solidify the performance of our method, we compare the performance of more recent methods. Based on fair comparisons, comparative experiments are performed without data augmentation and following our data preparation format (no extra input or pre-training). We show PSNR and the number of parameters in Table 2.

## 4.2 Results on Subtasks

In this section, we evaluate LW-ISP on subtasks to further explore the potential of end-to-end ISP model to reduce the calculation, comprehensively handle various tasks and generalize migration. Specifically, we demonstrate the effectiveness on denoising and enhancement.

Method	DnCNN[10]	TNRD[11]	BM3D[12]	WNNM[13]	KSVD[14]	EPLL[15]	CBDNet[16]	RIDNet[17]	VDN[18]	LW-ISP	MIRNet[60]
PSNR	23.66	24.73	25.65	25.78	26.88	27.11	30.78	38.71	39.28	<b>39.44</b>	39.72
SSIM	0.583	0.643	0.685	0.809	0.842	0.870	0.754	0.914	0.909	<b>0.918</b>	0.959

Table 4: Denoising results on the SIDD dataset [10]. Compared to the previous methods, our LW-ISP (numbers in bold) demonstrates a comparable performance. Blue font indicates the value higher than our method.

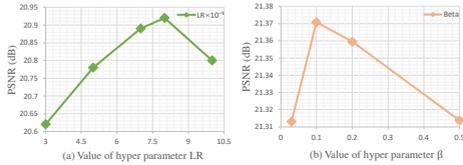


Figure 6: Sensitivity results of the parameters. (a) Learning Rate (LR): The initial learning rate in the training process. (b) Beta: The hyper-parameter  $\beta$  in Equation (1) for the distillation loss  $\mathcal{L}_d$ .

Backbone	FGAM	CCB	Distillation	PSNR (dB)
✓				20.92
✓	✓			21.17
✓		✓		21.28
✓	✓	✓		21.36
✓		✓	✓	21.40
✓	✓	✓	✓	<b>21.57</b>

Table 5: Ablation study of our method. Backbone refers to the basic architecture of Section 3.

**Image Denoising.** We train our network only on the training set of SIDD and directly evaluate it on the test images of both SIDD and DND datasets. Quantitative comparisons in terms of PSNR and SSIM metrics are summarized in Table 4 for SIDD. These experimental results show the excellent performance of our LW-ISP under such lightweight conditions, which has surpassed traditional data-driven algorithms. Furthermore, it is worth noting that our method provides better results while RIDNet [17] uses additional training data, yet VIDNet [59] and MIRNet [60] are much larger than LW-ISP.

**Image Enhancement.** Without using any additional data and tricks, we achieve quite competitive results on the LoL [62]. The PSNR can reach 20.18 dB, surpassing previous methods such as CRM [58] and MF [18].

### 4.3 Ablation Study

**Study on Hyper-parameters.** In Fig. 6 we show a sensitivity analysis of the parameters, which are used for LW-ISP training on the ISP pipeline and its subtasks. It can be observed that the LW-ISP backbone achieves the effect that we reported (20.92 dB) when the initial learning rate is  $8 \times 10^{-5}$ . In subfigure (b), we show the sensitivity study of  $\beta$  during LW-ISP training (without structural loss  $\mathcal{L}_s$ ). It can be easily observed that our performance boost holds for different values of  $\beta$ . The value of  $\alpha$  has been set at 0.4.

**Study on Architectural Components.** Table 5 shows that FGAM and CCB can bring gains of 0.25 dB and 0.36 dB respectively, and the combination of the two can bring 0.44 dB gains. The baseline LW-ISP contains the base U-Net with a global feature vector and is trained with  $\mathcal{L}_r + \mathcal{L}_s$ , denoted as the backbone. Moreover, the implicit features from clean image learning by distillation are helpful and boost performance. *More sensitivity studies of the hyper-parameters and ablation studies of FGAM can be found in the supplementary.*

**Study on Distillation Algorithm.** When designing the distillation algorithm, a key insight is that we believe that the main structure (UNet) can be divided into image understanding and image reconstruction processes respectively according to the down/up sampling steps. (1) Our method can learn the reconstruction features corresponding to clean images under the

supervision of the teacher during the upsampling reconstruction process. (2) We selected the outputs of the last three upsampling blocks. When performing ablation experiments at clean feature locations, the outputs of the four down/up sampling blocks are denoted as  $down_i$  and  $up_i$ , respectively. It turns out that the closer the supervision to the RGB output location, the more performance gains (baseline: 21.28dB,  $down_1 + down_2 + down_3$ : 21.11dB,  $down_4 + up_1 + up_2$ : 21.33dB,  $up_2 + up_3 + up_4$ : 21.40dB).

**Study on FGAM.** As to Fine-Grained Attention Module (FGAM), we use an addition operation (+) instead of multiplication (\*) to generate channel attention and spatial attention. The reason is that we are surprised to find that the training process corresponding to the addition operation is more stable, and the performance is better (+21.17dB, \*21.05dB). To perform a fair comparison, we experiment with the attention mechanism on the backbone. Specifically, channel attention and spatial attention respectively achieved 21.10 dB (+0.18) and 21.04 dB (+0.12). Our experimental results demonstrate that the combination fashion of channel attention and spatial attention in the FGAM is pretty essential. Unlike MIRNet [6] or CBAM [66], the attention map output by each branch will be added to the input feature point-to-point, and then the two will be concatenated across channel dimension. It is observed that: (a) The multiplication of the attention map and the input feature that will lead to sub-optimal results (21.05dB). Note that the baseline and the results after adding FGAM are 20.97 dB and 20.17 dB, respectively. (b) It is unnecessary to add a convolutional layer with a kernel size of  $1 \times 1$  to reduce the number of channels. It will reach a result of 20.11 dB (+0.14 dB), still 0.06 dB from the optimal result. (c) The FGAM position should be placed after the downsample block. If the position is reversed, it only reaches 20.04 dB (+0.07 dB), which is 0.13 dB away from the optimal result.

## 5 Discussion and Future Work

**Deep Application of ISP.** With the further collaborative development of ISP and AI vision, we believe that the collaboration between RAW data processed by neural networks and subsequent DL-based tasks will be more worth looking forward to in both theory and application. For instance, there is no unified conclusion on the objective evaluation standard of the processed photos. The effect of the smart ISP on the subsequent tasks may be a brand new viewpoint. What’s more, smart ISPs can also integrate and learn from downstream tasks such as image recognition. The cooperation between the CNN accelerator and the ISP also requires a unique design. The pipeline that derives from this integration of hardware-friendly tradition and smart ISP will be a new direction.

## 6 Conclusion

In this paper, we propose the LW-ISP to achieve real-time and high performance processing in smart ISP. The entire network implicitly learns the image mapping from RAW data to RGB photos, and utilizes the fine-grained attention module (FGAM), the contextual complement block (CCB) and heterogeneous distillation algorithm to reconstruct high-quality images. Abundant experiments show that LW-ISP has achieved state-of-the-art performance. The model parameters have been reduced by  $23 \times$  and the inference time has been accelerated by at least  $15 \times$ . Specifically, the results show that a U-Net (trained in the right manner) can replace much larger nets.

## References

- [1] Abdelrahman Abdelhamed, Stephen Lin, and Michael S Brown. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1692–1700, 2018.
- [2] Michal Aharon, Michael Elad, and Alfred Bruckstein. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on signal processing*, 54(11):4311–4322, 2006.
- [3] Saeed Anwar and Nick Barnes. Real image denoising with feature attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3155–3164, 2019.
- [4] Guillem Brasó and Laura Leal-Taixé. Learning a neural solver for multiple object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6247–6257, 2020.
- [5] 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.
- [6] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [7] Ashok Basur Chandrappa, Pradeep Kumar Nagaraj, Srikanth Vasudevan, Anantheswar Yelampalli Nagaraj, Krithika Jagadish, and Ankit Shah. Use of selfie sticks and iphones to record operative photos and videos in plastic surgery. *Indian journal of plastic surgery: official publication of the Association of Plastic Surgeons of India*, 50(1):82, 2017.
- [8] Dongliang Chang, Yifeng Ding, Jiyang Xie, Ayan Kumar Bhunia, Xiaoxu Li, Zhanyu Ma, Ming Wu, Jun Guo, and Yi-Zhe Song. The devil is in the channels: Mutual-channel loss for fine-grained image classification. *IEEE Transactions on Image Processing*, 29: 4683–4695, 2020.
- [9] Prashant Chaudhari, Franziska Schirmacher, Andreas Maier, Christian Riess, and Thomas Köhler. Merging-isp: Multi-exposure high dynamic range image signal processing. *arXiv preprint arXiv:1911.04762*, 2019.
- [10] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. *arXiv preprint arXiv:2204.04676*, 2022.
- [11] Yunjin Chen, Wei Yu, and Thomas Pock. On learning optimized reaction diffusion processes for effective image restoration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5261–5269, 2015.

- [12] Shen Cheng, Yuzhi Wang, Haibin Huang, Donghao Liu, Haoqiang Fan, and Shuaicheng Liu. Nbnnet: Noise basis learning for image denoising with subspace projection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4896–4906, 2021.
- [13] Yijia Cheng, Huanjing Yue, and Yan Mao. A lightweight convolutional neural network for camera isp. In *2021 IEEE 21st International Conference on Communication Technology (ICCT)*, pages 1346–1350. IEEE, 2021.
- [14] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007.
- [15] Linhui Dai, Xiaohong Liu, Chengqi Li, and Jun Chen. Awnet: Attentive wavelet network for image isp. In *European Conference on Computer Vision*, pages 185–201. Springer, 2020.
- [16] Chao Dong, Chen Change Loy, Kaifeng He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015.
- [17] Adriano Faggiani, Enrico Gregori, Luciano Lenzini, Valerio Luconi, and Alessio Vecchio. Smartphone-based crowdsourcing for network monitoring: opportunities, challenges, and a case study. *IEEE Communications Magazine*, 52(1):106–113, 2014.
- [18] Xueyang Fu, Delu Zeng, Yue Huang, Xiao-Ping Zhang, and Xinghao Ding. A weighted variational model for simultaneous reflectance and illumination estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2782–2790, 2016.
- [19] Shuhang Gu, Lei Zhang, Wangmeng Zuo, and Xiangchu Feng. Weighted nuclear norm minimization with application to image denoising. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2862–2869, 2014.
- [20] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*, 2020.
- [21] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward convolutional blind denoising of real photographs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1712–1722, 2019.
- [22] Wei Han, Zhengdong Zhang, Yu Zhang, Jiahui Yu, Chung-Cheng Chiu, James Qin, Anmol Gulati, Ruoming Pang, and Yonghui Wu. Contextnet: Improving convolutional neural networks for automatic speech recognition with global context. *arXiv preprint arXiv:2005.03191*, 2020.
- [23] Felix Heide, Markus Steinberger, Yun-Ta Tsai, Mushfiqur Rouf, Dawid Pająk, Dikpal Reddy, Orazio Gallo, Jing Liu, Wolfgang Heidrich, Karen Egiazarian, et al. Flexisp: A flexible camera image processing framework. *ACM Transactions on Graphics (TOG)*, 33(6):1–13, 2014.

- [24] Ming Hong, Yuan Xie, Cuihua Li, and Yanyun Qu. Distilling image dehazing with heterogeneous task imitation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3462–3471, 2020.
- [25] Yuenan Hou, Zheng Ma, Chunxiao Liu, and Chen Change Loy. Learning to steer by mimicking features from heterogeneous auxiliary networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8433–8440, 2019.
- [26] Ming-Chun Hsyu, Chih-Wei Liu, Chao-Hung Chen, Chao-Wei Chen, and Wen-Chia Tsai. Csanet: High speed channel spatial attention network for mobile isp. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2486–2493, 2021.
- [27] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.
- [28] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Dslr-quality photos on mobile devices with deep convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3277–3285, 2017.
- [29] Andrey Ignatov, Radu Timofte, Zhilu Zhang, Ming Liu, Haolin Wang, Wangmeng Zuo, Jiawei Zhang, Ruimao Zhang, Zhanglin Peng, Sijie Ren, et al. Aim 2020 challenge on learned image signal processing pipeline. In *European Conference on Computer Vision*, pages 152–170. Springer, 2020.
- [30] Andrey Ignatov, Luc Van Gool, and Radu Timofte. Replacing mobile camera isp with a single deep learning model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 536–537, 2020.
- [31] Andrey Ignatov, Cheng-Ming Chiang, Hsien-Kai Kuo, Anastasia Sycheva, and Radu Timofte. Learned smartphone isp on mobile npus with deep learning, mobile ai 2021 challenge: Report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2503–2514, 2021.
- [32] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [33] Byung-Hoon Kim, Joonyoung Song, Jong Chul Ye, and JaeHyun Baek. Pynet-ca: enhanced pynet with channel attention for end-to-end mobile image signal processing. In *European Conference on Computer Vision*, pages 202–212. Springer, 2020.
- [34] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.

- [35] Zhetong Liang, Jianrui Cai, Zisheng Cao, and Lei Zhang. Cameranet: A two-stage framework for effective camera isp learning. *IEEE Transactions on Image Processing*, 30:2248–2262, 2021.
- [36] Lin Liu, Xu Jia, Jianzhuang Liu, and Qi Tian. Joint demosaicing and denoising with self guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2240–2249, 2020.
- [37] Kangfu Mei, Juncheng Li, Jiajie Zhang, Haoyu Wu, Jie Li, and Rui Huang. Higher-resolution network for image demosaicing and enhancing. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3441–3448. IEEE, 2019.
- [38] Yiqun Mei, Yuchen Fan, and Yuqian Zhou. Image super-resolution with non-local sparse attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3517–3526, 2021.
- [39] Shervin Minaee, Yuri Boykov, Fatih Porikli, Antonio Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image segmentation using deep learning: A survey. *arXiv preprint arXiv:2001.05566*, 2020.
- [40] T Nishimura, Y Inoue, K Sugahara, S Kusunoki, T Kumamoto, S Nakagawa, M Nakaya, Y Horiba, and Y Akasaka. Three dimensional ic for high performance image signal processor. In *1987 International Electron Devices Meeting*, pages 111–114. IEEE, 1987.
- [41] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2337–2346, 2019.
- [42] Tobias Plotz and Stefan Roth. Benchmarking denoising algorithms with real photographs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1586–1595, 2017.
- [43] Daniel Wirzberger Raimundo, Andrey Ignatov, and Radu Timofte. Lan: Lightweight attention-based network for raw-to-rgb smartphone image processing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 808–816, 2022.
- [44] Rajeev Ramanath, Wesley E Snyder, Youngjun Yoo, and Mark S Drew. Color image processing pipeline. *IEEE Signal Processing Magazine*, 22(1):34–43, 2005.
- [45] Sivalogeswaran Ratnasingam. Deep camera: A fully convolutional neural network for image signal processing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [46] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [47] Eli Schwartz, Raja Giryes, and Alex M Bronstein. Deepisp: Toward learning an end-to-end image processing pipeline. *IEEE Transactions on Image Processing*, 28(2): 912–923, 2018.

- [48] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016.
- [49] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficientdet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10781–10790, 2020.
- [50] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018.
- [51] Jessica Van Brummelen, Marie O’Brien, Dominique Gruyer, and Homayoun Najjaran. Autonomous vehicle perception: The technology of today and tomorrow. *Transportation research part C: emerging technologies*, 89:384–406, 2018.
- [52] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2017.
- [53] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. Ieee, 2003.
- [54] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. *arXiv preprint arXiv:1808.04560*, 2018.
- [55] Kaixuan Wei, Ying Fu, Jiaolong Yang, and Hua Huang. 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*, pages 2758–2767, 2020.
- [56] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.
- [57] Chyuan-Tyng Wu, Leo F Isikdogan, Sushma Rao, Bhavin Nayak, Timo Gerasimow, Aleksandar Sutic, Liron Ain-kedem, 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.
- [58] Zhenqiang Ying, Ge Li, Yurui Ren, Ronggang Wang, and Wenmin Wang. A new image contrast enhancement algorithm using exposure fusion framework. In *International Conference on Computer Analysis of Images and Patterns*, pages 36–46. Springer, 2017.
- [59] Zongsheng Yue, Hongwei Yong, Qian Zhao, Lei Zhang, and Deyu Meng. Variational denoising network: Toward blind noise modeling and removal. *arXiv preprint arXiv:1908.11314*, 2019.

- [60] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for real image restoration and enhancement. *arXiv preprint arXiv:2003.06792*, 2020.
- [61] Syed Waqas Zamir, Aditya Arora, Salman Khan, Fahad Shahbaz Khan, and Ling Shao. Learning digital camera pipeline for extreme low-light imaging. *Neurocomputing*, 452: 37–47, 2021.
- [62] Bangcheng Zhan, Feng Li, and Ming Lu. Hdr synthesis technology for spaceborne cmos cameras based on virtual digital tdi. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 13:3824–3833, 2020.
- [63] Hongguang Zhang, Yuchao Dai, Hongdong Li, and Piotr Koniusz. Deep stacked hierarchical multi-patch network for image deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5978–5986, 2019.
- [64] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017.
- [65] Yu Zhu, Zhenyu Guo, Tian Liang, Xiangyu He, Chenghua Li, Cong Leng, Bo Jiang, Yifan Zhang, and Jian Cheng. Eednet: enhanced encoder-decoder network for autoisp. In *European Conference on Computer Vision*, pages 171–184. Springer, 2020.
- [66] Daniel Zoran and Yair Weiss. From learning models of natural image patches to whole image restoration. In *2011 International Conference on Computer Vision*, pages 479–486. IEEE, 2011.