

Learning a Reinforced Agent for Flexible Exposure Bracketing Selection

Zhouxia Wang^{1,2}, Jiawei Zhang^{1*}, Mude Lin¹, Jiong Wang¹, Ping Luo², Jimmy Ren¹

¹SenseTime Research, ²The University of Hong Kong

{zhouxi1212, zhjw1988, jimmy.sj.ren}@gmail.com, {linmude, wangjiong}@sensetime.com, pluo@cs.hku.hk

Abstract

Automatically selecting exposure bracketing (images exposed differently) is important to obtain a high dynamic range image by using multi-exposure fusion. Unlike previous methods that have many restrictions such as requiring camera response function, sensor noise model, and a stream of preview images with different exposures (not accessible in some scenarios e.g. some mobile applications), we propose a novel deep neural network to automatically select exposure bracketing, named *EBSNet*, which is sufficiently flexible without having the above restrictions. *EBSNet* is formulated as a reinforced agent that is trained by maximizing rewards provided by a multi-exposure fusion network (*MEFNet*). By utilizing the illumination and semantic information extracted from just a single auto-exposure preview image, *EBSNet* can select an optimal exposure bracketing for multi-exposure fusion. *EBSNet* and *MEFNet* can be jointly trained to produce favorable results against recent state-of-the-art approaches. To facilitate future research, we provide a new benchmark dataset for multi-exposure selection and fusion. Our code and proposed benchmark dataset will be released in <https://github.com/wzhouxiff/EBSNetMEFNet.git>

1. Introduction

Many real-world scenes have a very large dynamic range. However, the digital cameras, especially those on mobile devices, have limited dynamic range making them infeasible to capture a wide dynamic range image with only a single shot. Stack-based high dynamic range (HDR) imaging is a widely used technique to solve this issue by fusing several low dynamic range (LDR) images captured at different exposures [4, 5, 14, 17, 19, 20, 24, 25, 28]. For example, Debevec et al. [5] propose to use geometrically different exposure settings. However, this approach leads to a large exposure bracketing that has both long acquisi-




Figure 1. (a) is an AE preview image while (b) and (c) are the HDR images generated with the exposure bracketing selection methods proposed by Pourreza-Shahri et al. [23] and us. Both methods implement exposure bracketing selection only based on the AE preview image. Our proposed method can attain more details in saturated areas than [23]. Since Pourreza-Shahri et al. [23] focus on the illumination of the AE image, it does not know what the true brightness is in the saturated areas. However, our method considers the illumination and semantic information of the AE scene at the same time. Though the area is saturated, it is still possible to be recognized as billboard with semantic information. Therefore, our method tends to choose the exposure bracketing ($\{z_1, z_4, z_6\}$) that captured the details of the billboard.

tion and post-processing time, high memory requirement, as well as ghost artifacts yielded by motion. Therefore, selecting a small exposure bracketing is a key and important problem.

Most of the existing exposure bracketing selection methods propose to either find a minimal-bracketing set to cover the whole dynamic range of the scene [2] or to maximize the signal-to-noise ratio (SNR) [10, 7, 27, 30]. Although these methods have a decent performance, they have many limitations. For example, Hasinoff et al. [10] and Seshadri-nathan et al. [27] can only handle raw images that are linearly related to the scene irradiance. To be more flexible with both linear and non-linear operations on the camera e.g. the image signal processor (ISP), Barakat et al. [2], Gallo et al. [7] and Beek et al. [30] pre-estimate the camera response function based on [5]. However, all these methods need to employ either the dynamic range [2, 10] or the full irradiance histogram [7, 27, 30] of the scene by using a stream of preview images. Even though the preview images

*Corresponding author is Jiawei Zhang (Email: zhjw1988@gmail.com). This work was partially supported by HKU Seed Fund for Basic Research, Start-up Fund and Research Donation from SenseTime.

are easy to access sometimes, it is not user-friendly to preview the scene with different exposures. To alleviate this issue, Huang et al. [11] and Pourreza-Shahri et al. [23, 22] select exposures according to the scene information from the auto-exposure (AE) preview image and the camera parameters. As a result, they fail to capture information beyond the dynamic range from this single preview image, leading to sub-optimal performance.

Although neural networks are widely used to generate HDR image [6, 35, 31, 33, 24, 4], they have not been applied to select exposure bracketing. The proposed exposure bracketing selection network (EBSNet) can extract both illumination as well as semantic information from a single low-resolution LDR image. Unlike [11, 23, 22] only use illumination information from a single LDR image, the proposed EBSNet can utilize the information beyond the dynamic range of the single LDR image, such as the saturated light and the indistinguishable details in the dark grassland by considering the semantic information. Figure 1 shows an example whose input image is over-exposure in the billboard and the method proposed by Pourreza-Shahri et al. [23] fails to predict an exposure bracketing to cover all the details in these areas. However, with the analysis of semantic information and illumination distribution of the preview image, our model succeeds in finding the exposure that can capture the content under the billboard which is always informative. In addition, for some scenarios such as mobile applications where the raw data, the camera response function or the sensor noise model parameters are not accessible, the existing exposure bracketing selection methods [2, 10, 7, 27, 11, 23, 22, 30] that strongly rely on the above additional information are infeasible. However, the proposed EBSNet considers less of these information and can automatically learn to select out the appropriate exposure bracketing.

We employ Reinforcement Learning [29] to update the parameters in EBSNet. In this paper, we propose to train the proposed EBSNet as an agent, which is rewarded by the output of the multi-exposure fusion against the ground truth HDR images. Specifically, if the currently selected exposure bracketing generates a higher quality HDR image, the EBSNet will get a positive reward, and vice versa. We propose a multi-exposure fusion network (MEFNet) to merge the images of different exposures predicted by EBSNet and generate an HDR image. In the meantime, it will feedback a reward to EBSNet for the update of EBSNet. EBSNet and MEFNet can be trained jointly and the entire framework can produce favorable results against recent state-of-the-art methods.

To facilitate future research, we also collect a dataset. Our dataset contains 800 samples of various scenes, such as indoor and outdoor scenes, night sight and day sight scenes. For each scene, it consists of an image obtained under auto-exposure and ten images whose exposure time are related to the auto-exposure image. Specifically, the image captured

by using auto-exposure is used to determine the optimal exposure bracketing, while the remaining ten images with different exposures are the candidate images for constructing the predicted optimal exposure bracketing. These ten images are also used to generate the ground truth for HDR image generation.

The **contributions** of this work are as follows.

- The proposed EBSNet can use the illumination and semantic information of a low-resolution auto-exposure image to predict the optimal exposure bracketing without knowing the camera response function as well as the sensor noise model parameters and can also adapt to the images non-linear to the scene irradiance.
- The EBSNet is trained as an agent in reinforcement learning and rewarded with a novel MEFNet. The EBSNet and MEFNet can be trained jointly and they can perform favorably against state-of-the-art exposure bracketing selection methods.
- We provide a new benchmark dataset to facilitate the future studies of exposure bracketing selection.

2. Related Work

Exposure Bracketing Selection. Most of the existing Exposure Bracket Selection methods are based on the estimation of the dynamic range. Barakat et al. [2] tend to find the minimal-bracketing sets for attaining HDR images by assuming a linear camera response function and a uniform irradiance distribution. Hasinoff et al. [10] focus on achieving higher worst-case SNR by minimizing the number of exposures or the total exposure duration. Grossberg et al. [9] optimize the exposure bracket by minimizing an objective function that is based on the derivative of the response function. While the above three methods are based on the estimation of the extent of the dynamic range, Gallo et al. [7], Seshadrinathan et al. [27] and Beek et al. [30] present a method to estimate the full irradiance histogram of the scene by capturing a stream of images and a strategy to select the set of exposures that need to be acquired. Also, some methods do not require the scene irradiance to be known [11, 22, 23] and they use the scene information from the normal or auto-exposure and the camera parameters to perform exposure bracketing.

For releasing our model from the estimation of the dynamic range and response function, our proposed model only has a view on the auto-exposure low-resolution preview LDR image. Different from the similar works [11, 22, 23] who just focus on the illumination of AE image, our method also take consideration on the semantic information of the preview image by a neural network which can utilize the information beyond the dynamic range of the AE image, such as the saturated light and the indistinguishable details in the dark grassland.

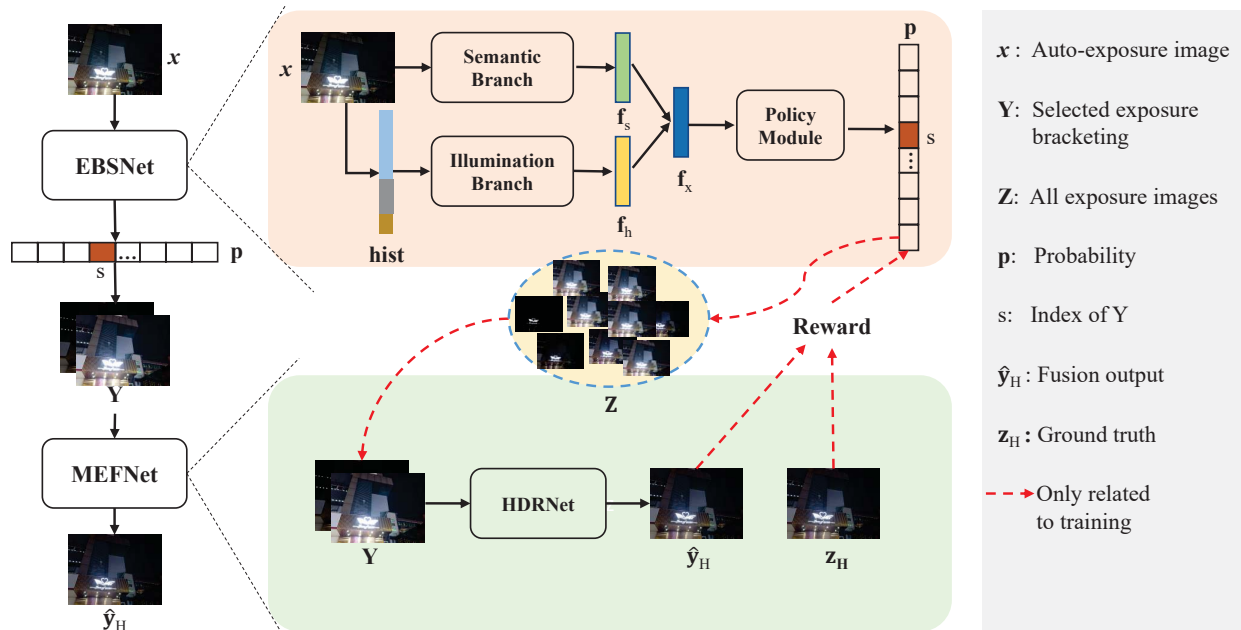


Figure 2. This figure shows the pipeline of our proposed method. First, it captures an auto-exposure low-resolution preview image, and then feeds it into an Exposure Bracketing Selection Network (EBSNet) to determine an optimal exposure bracketing according to the extracted illumination and semantic information. Multi-Exposure Fusion Network (MEFNet) will take the selected exposure bracketing as input and generate an HDR image. When training EBSNet, We regard it as an agent and it gets the reward from MEFNet. EBSNet and MEFNet can be trained jointly. Please see the manuscript for more details.

Multi-Exposure Fusion. Traditional Multi-Exposure Fusion methods mainly use hand-crafted features. Mertens et al. [20] fuse the images by using simple quality measures such as saturation and contrast. Ma et al. [19] decompose the image patch into three components and fuse them separately to obtain the desired one. There are also some methods of attaining fusion weights by using image filters. For example, Li et al. [17] decompose an image into two layers and employ a guided filtering-based weighted average technique to make full use of the spatial consistency for exposure fusion. Raman et al. [25] use bilateral filters to preserve details in both bright and dark regions of the scene. As the previous methods neglect the semantic information of the images, [4, 24] introduce CNNs into Multi-Exposure Fusion to extract the high-level features with generated ground truth [4] or without reference [24]. Our MEFNet aims to provide a reward for the update of the EBSNet. It is easy to replace our MEFNet with the above MEF method. However, for modeling the interaction between the exposure bracketing selection and exposure fusion and making them can be trained jointly, we tend to implement exposure fusion with a neural network.

Reinforcement Learning. Recently, reinforcement learning (RL) has been applied to all kind of computer vision tasks, such as classification [1, 21], detection [12], visual captioning [18, 16, 26], machine translation [32], and so on. There are also some exposure-related works up-

dated via reinforcement learning [34, 36]. However, Yang et al. [34] aim to build a system that users can adapt the exposure of a single image personally via giving a score to a preview image and Yu et al. [36] use RL to learn local exposures for exposure fusion with segmentation and carefully retouch. Our model tends to use RL to determine the optimal exposure bracketing for the generation of an HDR image by analyzing the preview image and get a reward directly from a multi-exposure fusion network. By joint training, our EBSNet and MEFNet can get benefit from each other.

3. Method

The purpose of this method is to select an optimal exposure bracketing for HDR image generation and it contains two components: EBSNet and MEFNet which are showed in Figure 2. Unlike most of the existing methods that need to consider a sequence of preview images with different exposures, the proposed EBSNet only needs to analyze a single AE low-resolution preview image by considering both the illumination distribution and the semantic information. Since we cannot directly get the ground truth of the exposure bracketing, we propose to update EBSNet via RL [29] guided by a reward generated from exposure fusion and its HDR ground truth. Even though any exposure fusion method can be used in RL framework to get the reward, we propose a network MEFNet to fuse the selected expo-

sure bracketing since both EBSNet and MEFNet can benefit from each other through joint training.

3.1. Exposure Bracketing Selection Network (EBSNet)

The input of the proposed EBSNet is a low-resolution AE LDR preview image \mathbf{x} and EBSNet predicts a bracketing of K different exposures $\mathbf{Y} = \{\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{K-1}\}$ which may cover the dynamic range of the scene as large as possible. Since there is only one LDR image \mathbf{x} , it is difficult to select out an optimal exposure bracketing only based on the illumination distribution of \mathbf{x} especially in the area where is over-exposure or extremely dark. For a better selection, EBSNet takes the semantic information of the preview image into consideration and it is favorable. As Figure 1 shown, though the billboard in the AE image is saturated and its illumination information has been discarded, it is still possible to be recognized as a billboard with its semantic information. Since billboards are used for displaying informative content, our model tends to get an exposure bracketing that captured the details of the billboard. Therefore, the proposed EBSNet consists of two branches. One is **semantic branch** for semantic feature extraction while the other is **illumination branch** for illumination feature extraction (as shown in Figure 2).

Semantic Branch The semantic feature extraction is constructed as Alexnet [15]. It takes the preview image \mathbf{x} as input and extracts the semantic feature \mathbf{f}_s from the first fully-connected layer. We can formulate it as:

$$\mathbf{f}_s = \text{Alexnet}(\mathbf{x}). \quad (1)$$

Illumination Branch In order to utilize the illumination distribution globally and locally, we estimate the histogram of \mathbf{x} with a 3-levels spatial pyramid. For u -th level, we divide \mathbf{x} into $2^{(u-1)} \times 2^{(u-1)}$ non-overlapped patches and a histogram is estimated from every patch with 32 bins. By concatenating all the histograms, we get feature **hist** and send it to a neural network that consists of two convolutional layers and one fully-connected layer. Finally, we attain the illumination feature \mathbf{f}_h . By denoting this neural network as Histnet, we can represent \mathbf{f}_h as:

$$\mathbf{f}_h = \text{Histnet}(\mathbf{hist}). \quad (2)$$

To fuse these two features, we introduce an additional fully-connected layer which takes the concatenation of \mathbf{f}_s and \mathbf{f}_h as inputs and gets the final scene feature \mathbf{f}_x of preview image \mathbf{x} which is shown as:

$$\mathbf{f}_x = \text{fc}([\mathbf{f}_s, \mathbf{f}_h]). \quad (3)$$

Policy Module Similar to the existing methods, e.g. [2, 10], the exposure \mathbf{y}_k ($k \in [0, K-1]$) in the selected optimal exposure bracketing \mathbf{Y} is included in a large exposure

bracketing $\mathbf{Z} = \{\mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_{J-1}\}$ which has J different exposures and contains a large dynamic range. In this way, there are totally $N=C_J^K$ possible exposure bracketing in a scene and the candidate exposure bracketing is denoted as $\mathbf{A} = \{\mathbf{Y}_0, \mathbf{Y}_1, \dots, \mathbf{Y}_{N-1}\}$.

Given the scene feature \mathbf{f}_x , a policy module is used to predict a probability distribution \mathbf{p} over the candidate exposure bracketing \mathbf{A} and the exposure bracketing \mathbf{Y}_s with the highest probability will be sent into MEFNet to further generate the HDR image. The policy module can be described as follows:

$$\begin{aligned} \mathbf{p} &= \text{Softmax}(\mathbf{W}\mathbf{f}_x + \mathbf{b}) \\ s &= \text{argmax}(\mathbf{p}), \end{aligned} \quad (4)$$

where \mathbf{W} and \mathbf{b} are the weights and bias of a fully-connected layer, and s is the index of highest probability in \mathbf{p} .

3.2. Multi-Exposure Fusion Network (MEFNet)

By considering the K LDR images in \mathbf{Y}_s selected by EBSNet, exposure fusion can generate an HDR image $\hat{\mathbf{y}}_H$. Also, exposure fusion is used to update EBSNet. Since there is no ground truth for the selection procedure, we propose to use RL scheme for the training of EBSNet. With RL, the agent (EBSNet) will decide the action (one of the exposure bracketings in \mathbf{A}) in an environment (multi-exposure fusion) to maximize the cumulative reward according to the current state (quality of generated HDR image).

Although a differentiable environment is not necessary for RL framework, we still adopt a learnable neural network MEFNet derived from the HDRNet [8] for multi-exposure fusion. In this way, both EBSNet and MEFNet can benefit from each other through joint training. MEFNet can be expressed as:

$$\hat{\mathbf{y}}_H = \text{HDRNet}(\mathbf{Y}_s). \quad (5)$$

3.3. Learning

MEFNet loss When training MEFNet, it minimizes the Charbonnier Loss [3] which is defined as:

$$\mathcal{L}_f(\hat{\mathbf{y}}_H, \mathbf{z}_H) = \sum_i \sum_j \sqrt{(\hat{\mathbf{y}}_H(i, j) - \mathbf{z}_H(i, j))^2 + \epsilon^2}, \quad (6)$$

where (i, j) is the position in the image and $\epsilon = 10^{-3}$. To generate the ground truth HDR image \mathbf{z}_H , we fuse all images in \mathbf{Z} with an existing exposure fusion method [20].

EBSNet Loss In RL, the agent needs to gather as much reward as possible. In this paper, the reward is defined as the difference between the current and previous peak signal-to-noise ratio (PSNR) of the HDR image as:

$$\mathcal{R}(\hat{\mathbf{y}}_H^{v-1}, \hat{\mathbf{y}}_H^v, \mathbf{z}_H) = \text{PSNR}(\hat{\mathbf{y}}_H^v, \mathbf{z}_H) - \text{PSNR}(\hat{\mathbf{y}}_H^{v-1}, \mathbf{z}_H), \quad (7)$$




Figure 3. One sample in our dataset. It consists of an auto-exposure image \mathbf{x} and ten images $\{\mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_9\}$ whose exposure time are related to the auto-exposure. The exposure durations are also shown here. Best viewed in color; zoom in for additional details.

in which $\hat{\mathbf{y}}_H^v$ is the estimated HDR image generated from the selected exposure bracketing in v -th step. To update EBSNet, we need to minimize the following loss function as:

$$\mathcal{L} = -\mathbb{E}_s[\mathcal{R}(\hat{\mathbf{y}}_H^{v-1}, \hat{\mathbf{y}}_H^v, \mathbf{z}_H)] \quad (8)$$

and

$$\nabla \mathcal{L} = -\mathbb{E}_s[\nabla \log \mathbf{p}(s) \mathcal{R}((\hat{\mathbf{y}}_H^{v-1}, \hat{\mathbf{y}}_H^v, \mathbf{z}_H))], \quad (9)$$

where s is the action chosen by EBSNet represented a type of exposure bracketing and $\mathbb{E}_s[\cdot]$ is the expectation over s .

3.4. Training Procedure

Since the updating of EBSNet is determined by the reward feedbacked from the MEFNet and the learning of MEFNet is affected by the selection of EBSNet, an appropriate weight initialization of both networks is necessary. Therefore, before jointly training these two networks, we train them separately to get pretrain models. The training procedure includes three steps. First, we train MEFNet with K images that randomly selected from the candidate exposure bracketing \mathbf{A} for a more reliable reward to update EBSNet. Second, we use the trained MEFNet to generate HDR images and calculate the PSNRs w.r.t. the ground truth HDR image. The candidate exposure bracketing that has the top 1 PSNR among all the candidate exposure bracketings is set as the initial target action of the scene and is used for the coarse training of EBSNet. Third, we train EBSNet and MEFNet iteratively by every 10 epochs with the pretrained model attained in step one and step two. At that time, the input of MEFNet is selected by EBSNet, while the EBSNet is updated by the PSNR of the output of MEFNet and the ground truth HDR image.

4. Experiments and Analysis

4.1. Dataset

When training the proposed framework, an auto-exposure LDR preview image as well as a stream of images with different exposures, which cover a large dynamic

range, are needed. As there lacks a dataset that satisfies these requirements, we propose a new dataset that contains about 800 different scenes including day and night as well as indoor and outdoor. Specifically, each scene includes one auto-exposure preview image \mathbf{x} , $J = 10$ images \mathbf{Z} captured under different exposures, and one HDR image \mathbf{Z}_H generated by fusing all the ten images with an existing exposure fusion method [20]. The data are collected by a mobile hold on a tripod to avoid the motion between different exposures. The exposure time of the AE preview denotes as t_a and that of the rest ten images denote as $2t_a, 1.5t_a, t_a, 2^{-1}t_a, 2^{-2}t_a, 2^{-3}t_a, 2^{-4}t_a, 2^{-5}t_a, 2^{-6}t_a, 2^{-7}t_a$ respectively. The rest of the camera settings are fixed for each exposure in a scene. The corresponding images are named as $\{\mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_9\}$ and one of the samples is shown in Figure 3.

The dataset is randomly split into 3 parts: a training set contains 600 images, a validation set contains 100 images, and a test set contains 100 images.

4.2. Experiment Settings

Implementation details In this work, the length of each selected bracketing K is set to 3 as Barakat et al. [2] have shown that three images can capture the luminance of a scene adequately in most of the cases. Then, the total number of possible selected exposure bracketing is $N = C_{10}^3 = 120$. As some exposure bracketing are almost impossible to yield an acceptable HDR image, e.g. $\{\mathbf{z}_7, \mathbf{z}_8, \mathbf{z}_9\}$ which are too dark to capture any details in the dark regions, we only consider $N' = 36$ exposure bracketing in the candidate exposure bracketing \mathbf{A} in EBSNet. Please see the supplemental material for more details.

While training, the input size of EBSNet and MEFNet are 224×224 and 512×512 respectively. For data augmentation, we randomly crop, flip and rotate the images. The proposed network is trained with Adam [13]. The learning rates of MEFNet in the first step and EBSNet in the second step are set to 10^{-3} and 10^{-4} respectively. While in the third step, their learning rate is multiplied with 0.1. The batch size is 8.

Method	Barakat [2]	Beek [30]	Pourreza-Shahri [23]	EBSNet+EF	EBSNet+MEFNet
number	3.12	3.71	2.48	3.00	3.00

Table 1. Average number of selected exposure bracketing of each method. The number of images selected from [2] and [30] varies from scene to scene and the maximum number of images selected from [23] is 3. In the proposed framework, the number of the selected images K is fixed to 3. Please see Sec. 4.4 for more analysis of K of the proposed method in details.

Method	PSNR
Barakat [2]	27.31
Pourreza-Shahri [23]	26.44
Beek [30]	27.46
EBSNet+EF[20]	28.15
EBSNet+MEFNet	28.41

Table 2. Comparison of the state-of-the-art methods and ours in terms of PSNR. Both ‘EBSNet+EF’ and ‘EBSNet+MEFNet’ use the proposed exposure bracketing selection network. The only difference is that ‘EBSNet+EF’ uses exposure fusion proposed by [20] while ‘EBSNet+MEFNet’ uses the proposed multi-exposure fusion network. It shows that the proposed method can benefit from the joint training of EBSNet and MEFNet. In addition, the proposed method performs better than the existing methods.

4.3. Comparison with State-of-the-art Methods

To validate the effectiveness of our proposed method on exposure bracketing selection, we compare the performance of our method against the state-of-the-art approaches proposed by Barakat et al. [2], Pourreza-Shahri and Kehtarnavaz [23], and Beek et al. [30].

For a fair comparison, we use the exposure fusion method proposed by [20] to generate the HDR image with the selected exposure bracketing from the above methods as well as the proposed EBSNet (denoted as EBSNet+EF). And our whole framework with both the proposed exposure bracketing selection network as well as the multi-exposure fusion network is denoted as EBSNet+MEFNet. Besides, the average number of the selected exposure bracketing of these methods is close to three which is shown in Table 1.

According to Table 2, EBSNet+EF performs comparable relative to the existing methods. And EBSNet+MEFNet can further improve the performance through joint training which can also be seen from Figure 5. As shown in Figure 4, both the billboard at night and the building outside in the daytime are saturated in the AE preview x . The methods proposed by Barakat et al. [2] and Beek et al. [30] are based on the estimation of the dynamic range from a stream of preview images with different exposures. In this way, they can consider the histogram of the whole scene irradiance for selection and the saturated regions can be recovered well. As their objective is to estimate a minimal bracketing to achieve a worst-case SNR, they tend to choose the short exposure. As a result, the generated dark regions, e.g. the pink building, seem not as bright as the ground truth. Since Pourreza-Shahri et al. [23] only considers the illumination

Method	PSNR
Semantic	27.15
Illumination	27.40
EBSNet+MEFNet	28.41

Table 3. Effectiveness of semantic information according to PSNR. The proposed method contains both the semantic branch as well as the illumination branch. Even though only the illumination branch performs better than only the semantic branch, the proposed method can achieve better results by considering the semantic information.

K	1	2	3	10
PSNR	26.06	27.53	28.41	30.14
time/ms	218.43	301.21	352.27	811.25

Table 4. Performance and efficiency of different numbers of images K in the selected exposure bracketing in terms of PSNR and running time. The running time is about MEFNet in a mainstream Android device. The input size is 3000×4000 . The proposed method will generate better results with a larger number of images. Though the performance is higher when $K = 10$, it costs a lot of time while generating the HDR image. We choose to use $K = 3$ in the proposed method.

of the preview and does not know the brightness of the saturated region, it fails to recover these regions. By considering semantic information through training, the proposed method can better recover the saturated regions (the billboard in (e) and the building in (k)) and keep the brightness of dark regions (the pink building in (e)).

4.4. Ablation study

Illumination Distribution v.s. Semantic Information In order to validate the effectiveness of semantic information in EBSNet, we compare the proposed framework with two variant networks: removing the semantic branch and removing the illumination branch. Table 3 shows that the proposed method performs worse if either of the part is removed. The feature map in Figure 6 (b) shows that the network trained with semantic branch pays more attention to the lights as well as billboards and Figure 6 (d) can suppress the saturated regions better than Figure 6 (c) which is only based on illumination information.

Determination of K Even though Barakat et al. [2] have shown that three images can capture the luminance of a scene adequately in most cases, we also take some exper-

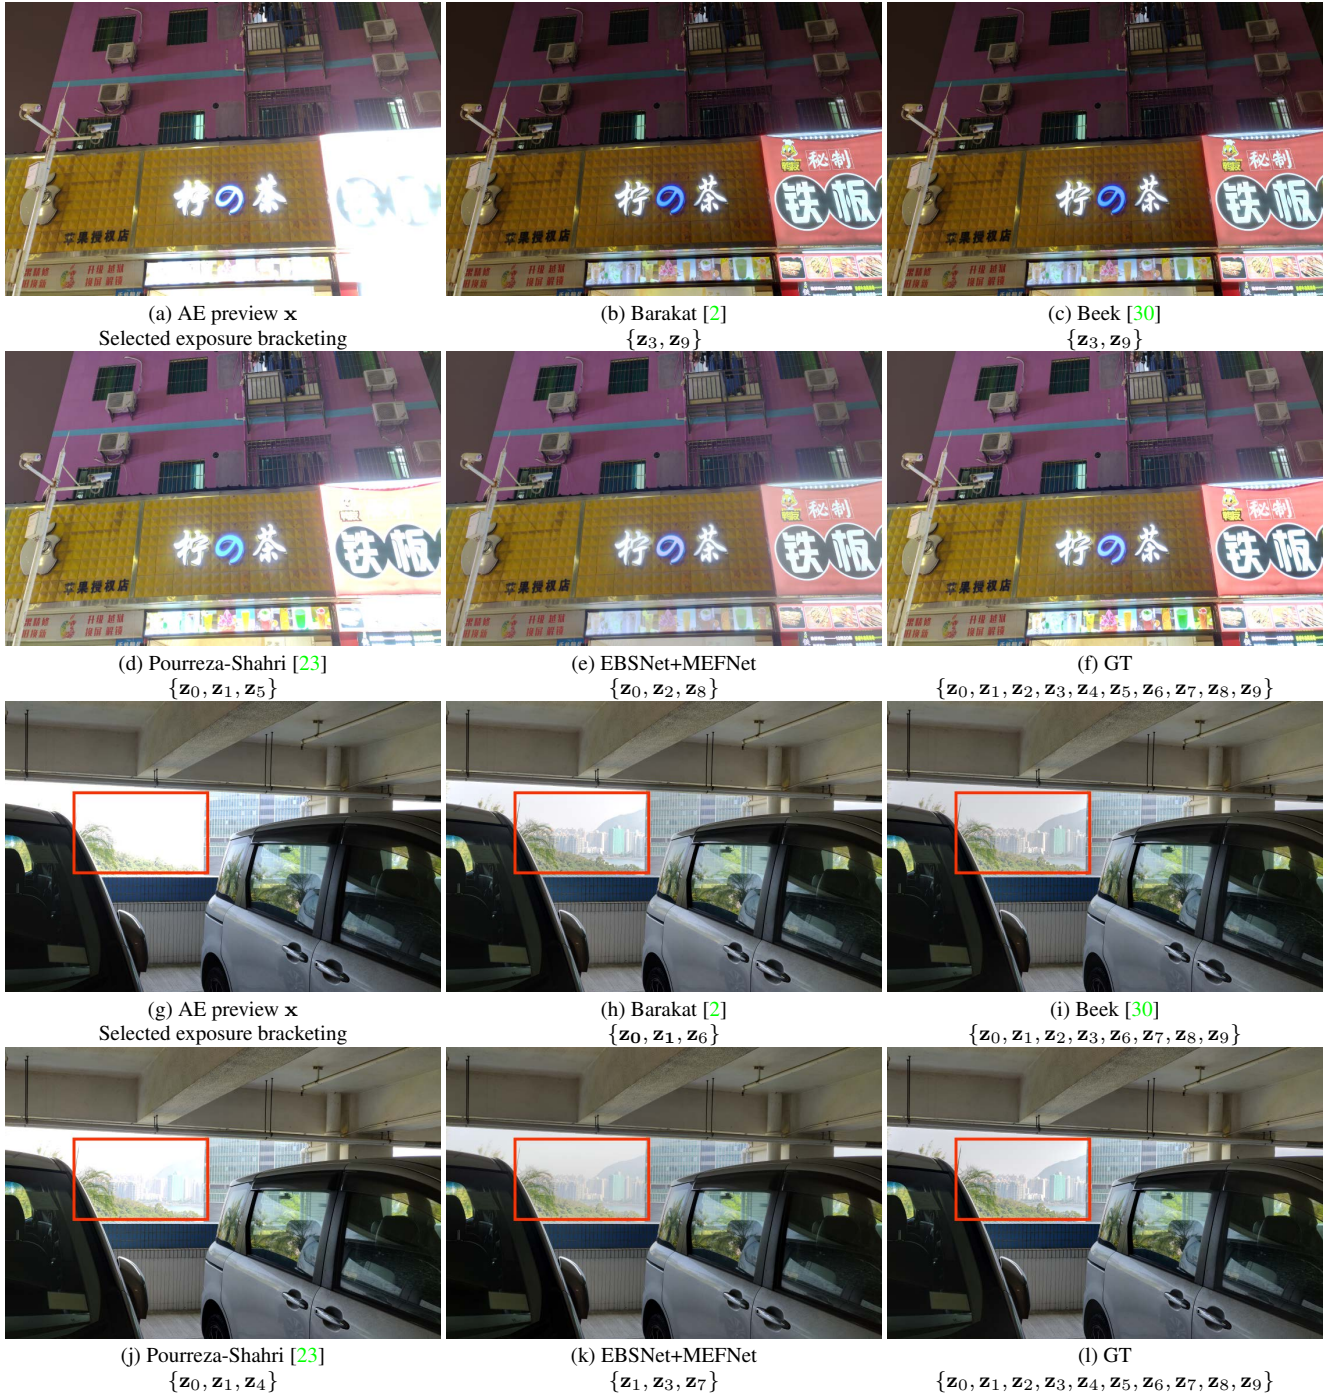


Figure 4. Comparison with the state-of-the-art exposure bracketing selection methods under night and daytime. By considering a stream of preview images with different exposures and the camera response function, Barakat et al. [2] and Beek et al. [30] can recover the saturated regions well, e.g. the billboard in (b)(c) and the building in (h)(i). However, their objective is to estimate a minimal bracketing to get a worst-case SNR, they fail to generate a bright enough building in (b) and (c). Both Pourreza-Shahri et al. [23] and the proposed method only use one AE preview image to select the exposure bracketing. As Pourreza-Shahri et al. [23] fail to consider the semantic information, it cannot recover the saturated regions, e.g. the billboard in (d) and the building in (j). With the help of semantic information, the proposed method can generate a better result both in the dark region in (e) and the saturated region in (k). Best viewed in color; zoom in for additional details.




Figure 5. Effectiveness of joint training. (a) is the AE preview image, (b) and (c) use [20] and the proposed MEFNet to generate HDR image with the exposure bracketing selected by the proposed EBSNet respectively, (d) is the ground truth. With jointly training EBSNet with MEFNet, MEFNet can suppress the saturated regions better than Mertens et al. [20] especially for the characters in the billboard since our objective is to make the fusion result of MEFNet as close as to the ground truth. Best viewed in color; zoom in for additional details.




Figure 6. Effectiveness of the semantic branch. (a) is the AE preview image, (b) is the feature map from the semantic branch of the proposed EBSNet, (c) and (d) are the generated HDR images from EBSNet without and with the semantic branch. The feature map from the semantic branch pays more attention to the saturated regions according to (b). As a result, (d) decides to select shorter exposures $\{z_1, z_4, z_7\}$ and can recover the saturated regions better than (c). Best viewed in color; zoom in for additional details.

iments to compare the different number of images K in the selected exposure bracketing. The results are shown in Table 4. When K equals to 1, it is a single image enhancement task like HDRNet [8]. It shows that the proposed framework will generate a more accurate HDR image with more LDR images as the input of MEFNet. When K equals to 10, the proposed MEFNet considers all the candidate exposures and generates the HDR images the most similar to those from the [20]. Though the performance is higher when $K = 10$, it is time-consuming. As shown in the table, the computational cost of MEFNet increases with more selected exposures, especially when we fuse the K images with full-resolution (3000×4000 of the proposal dataset). For the trade-off of performance and efficiency, we choose to use $K = 3$ in the proposed method. Besides, the running time of EBSNet is about 14ms with any K since it only influences the output size of the last fully-connected layer of EBSNet (C_{10}^K). We test our model in a mainstream Android

smartphone.

5. Conclusion

In this paper, we propose a neural network EBSNet to select exposure bracketing by extracting both the illumination as well as the semantic information from only one single AE preview image. It is free from knowing any additional information in prior which makes EBSNet flexible and can adapt to more mobile applications. In order to train EBSNet, we treat it as an agent in RL and reward it by another neural network, MEFNet, which is used for multi-exposure fusion. By jointly training EBSNet and MEFNet, the proposed exposure bracketing selection can perform favorably against state-of-the-art methods both quantitatively and qualitatively. We also provide a new dataset to facilitate future studies of exposure bracketing selection.

References

- [1] Jimmy Ba, Volodymyr Mnih, and Koray Kavukcuoglu. Multiple object recognition with visual attention. *arXiv preprint arXiv:1412.7755*, 2014. 3
- [2] Neil Barakat, A Nicholas Hone, and Thomas E Darcie. Minimal-bracketing sets for high-dynamic-range image capture. *TIP*, 2008. 1, 2, 4, 5, 6, 7
- [3] Jonathan T Barron. A general and adaptive robust loss function. In *CVPR*, 2019. 4
- [4] Jianrui Cai, Shuhang Gu, and Lei Zhang. Learning a deep single image contrast enhancer from multi-exposure images. *TIP*, 2018. 1, 2, 3
- [5] Paul E Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH*, 2008. 1
- [6] Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafat K Mantiuk, and Jonas Unger. Hdr image reconstruction from a single exposure using deep cnns. *TOG*, 2017. 2
- [7] Orazio Gallo, Marius Tico, Roberto Manduchi, Natasha Gelfand, and Kari Pulli. Metering for exposure stacks. In *Computer Graphics Forum*, 2012. 1, 2
- [8] Michaël Gharbi, Jiawen Chen, Jonathan T Barron, Samuel W Hasinoff, and Frédo Durand. Deep bilateral learning for real-time image enhancement. *TOG*, 2017. 4, 8
- [9] Michael D Grossberg and Shree K Nayar. High dynamic range from multiple images: Which exposures to combine. In *ICCVW*, 2003. 2
- [10] Samuel W Hasinoff, Frédo Durand, and William T Freeman. Noise-optimal capture for high dynamic range photography. In *CVPR*, 2010. 1, 2, 4
- [11] Kun-Fang Huang and Jui-Chiu Chiang. Intelligent exposure determination for high quality hdr image generation. In *ICIP*, 2013. 2
- [12] Zequn Jie, Xiaodan Liang, Jiashi Feng, Xiaojie Jin, Wen Lu, and Shuicheng Yan. Tree-structured reinforcement learning for sequential object localization. In *NIPS*, 2016. 3
- [13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv*, 2014. 5
- [14] Fei Kou, Zhengguo Li, Changyun Wen, and Weihai Chen. Multi-scale exposure fusion via gradient domain guided image filtering. In *ICME*, 2017. 1
- [15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 4
- [16] Lijun Li and Boqing Gong. End-to-end video captioning with multitask reinforcement learning. In *WACV*, 2019. 3
- [17] Shutao Li, Xudong Kang, and Jianwen Hu. Image fusion with guided filtering. *TIP*, 2013. 1, 3
- [18] Siqi Liu, Zhenhai Zhu, Ning Ye, Sergio Guadarrama, and Kevin Murphy. Improved image captioning via policy gradient optimization of spider. In *ICCV*, 2017. 3
- [19] Kede Ma, Hui Li, Hongwei Yong, Zhou Wang, Deyu Meng, and Lei Zhang. Robust multi-exposure image fusion: a structural patch decomposition approach. *TIP*, 2017. 1, 3
- [20] Tom Mertens, Jan Kautz, and Frank Van Reeth. Exposure fusion: A simple and practical alternative to high dynamic range photography. In *Computer graphics forum*, 2009. 1, 3, 4, 5, 6, 8
- [21] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *NIPS*, 2014. 3
- [22] Reza Pourreza-Shahri and Nasser Kehtarnavaz. Automatic exposure selection for high dynamic range photography. In *ICCE*, 2015. 2
- [23] Reza Pourreza-Shahri and Nasser Kehtarnavaz. Exposure bracketing via automatic exposure selection. In *ICIP*, 2015. 1, 2, 6, 7
- [24] K Ram Prabhakar, V Sai Srikar, and R Venkatesh Babu. Deepfuse: A deep unsupervised approach for exposure fusion with extreme exposure image pairs. In *ICCV*, 2017. 1, 2, 3
- [25] Shanmuganathan Raman and Subhasis Chaudhuri. Bilateral filter based compositing for variable exposure photography. In *Eurographics*, 2009. 1, 3
- [26] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. In *CVPR*, 2017. 3
- [27] Kalpana Seshadrinathan, Sung Hee Park, and Oscar Nestares. Noise and dynamic range optimal computational imaging. In *ICIP*, 2012. 1, 2
- [28] Rui Shen, Irene Cheng, Jianbo Shi, and Anup Basu. Generalized random walks for fusion of multi-exposure images. *TIP*, 2011. 1
- [29] Richard S Sutton, Andrew G Barto, et al. *Introduction to reinforcement learning*. MIT press Cambridge, 1998. 2, 3
- [30] Peter van Beek. Improved image selection for stack-based hdr imaging. *arXiv preprint arXiv:1806.07420*, 2018. 1, 2, 6, 7
- [31] Shangzhe Wu, Jiarui Xu, Yu-Wing Tai, and Chi-Keung Tang. Deep high dynamic range imaging with large foreground motions. In *ECCV*, 2018. 2
- [32] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv*, 2016. 3
- [33] Qingsen Yan, Dong Gong, Qinfeng Shi, Anton van den Hengel, Chunhua Shen, Ian Reid, and Yanning Zhang. Attention-guided network for ghost-free high dynamic range imaging. In *CVPR*, 2019. 2
- [34] Huan Yang, Baoyuan Wang, Noranart Vesdapunt, Minyi Guo, and Sing Bing Kang. Personalized exposure control using adaptive metering and reinforcement learning. *IEEE transactions on visualization and computer graphics*, 2018. 3
- [35] Xin Yang, Ke Xu, Yibing Song, Qiang Zhang, Xiaopeng Wei, and Rynson WH Lau. Image correction via deep reciprocating hdr transformation. In *CVPR*, 2018. 2
- [36] Runsheng Yu, Wenyu Liu, Yasen Zhang, Zhi Qu, Deli Zhao, and Bo Zhang. Deepexposure: Learning to expose photos with asynchronously reinforced adversarial learning. In *Advances in Neural Information Processing Systems*, pages 2149–2159, 2018. 3