

Semi-Supervised High Dynamic Range Image Reconstructing via Bi-Level Uncertain Area Masking

Wei Jiang*, Jiahao Cui*, Yizheng Wu, Zhan Peng, Zhiyu Pan[†], Zhiguo Cao

School of AIA, Huazhong University of Science and Technology
{wei_jiang,smartcjh,yzww21,peng_zhan,zhiyupan,zgcao}@hust.edu.cn

Abstract

Reconstructing high dynamic range (HDR) images from low dynamic range (LDR) bursts plays an essential role in the computational photography. Impressive progress has been achieved by learning-based algorithms which require LDR-HDR image pairs. However, these pairs are hard to obtain, which motivates researchers to delve into the problem of annotation-efficient HDR image reconstructing: how to achieve comparable performance with limited HDR ground truths (GTs). This work attempts to address this problem from the view of semi-supervised learning where a teacher model generates pseudo HDR GTs for the LDR samples without GTs and a student model learns from pseudo GTs. Nevertheless, the confirmation bias, *i.e.*, the student may learn from the artifacts in pseudo HDR GTs, presents an impediment. To remove this impediment, an uncertainty-based masking process is proposed to discard unreliable parts of pseudo GTs at both pixel and patch levels, then the student can learn from the trusted areas. With this novel masking process, our semi-supervised HDR reconstructing method not only outperforms previous annotation-efficient algorithms, but also achieves comparable performance with up-to-date fully-supervised methods by using only 6.7% HDR GTs.

Code — <https://github.com/JW20211/SmartHDR>

Introduction

Human eyes adjust the local exposure gain dynamically with retinal ganglion cells. Hence, views with high dynamic range (HDR) illumination can be presented in the brain (Larson 1998). Current digital imaging systems, where sensors have limited light sensitivity ranges (Hoefflinger 2007), can only reflect the absolute light intensity (Niu et al. 2021), which produces low dynamic range (LDR) images. For the goal of *what you see is what you get*, researchers try to fuse the differently exposed LDR images into HDR images (Kalantari, Ramamoorthi et al. 2017). However this pipeline causes ghosting artifacts for dynamic scenes.

To remove ghosting artifacts, most modern methods fall into a learning-based paradigm (Liu et al. 2021; Chung and

*These authors contributed equally.

[†]Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

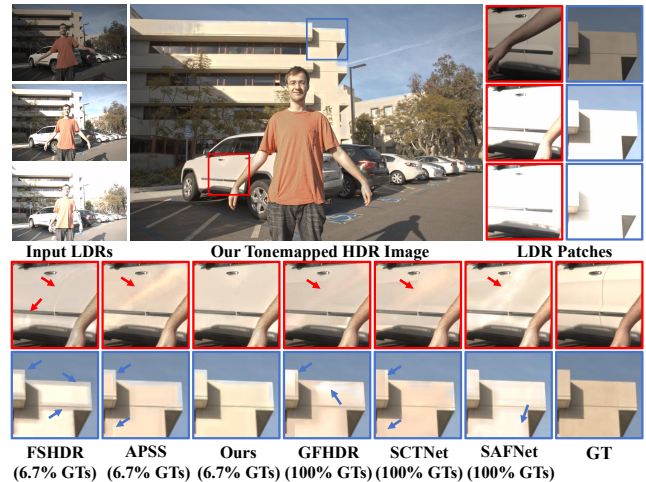


Figure 1. Qualitative comparison with prior arts. The illustrated example shows that, trained with only 6.7% GTs, our method outperforms previous annotation-efficient methods and achieve comparable or even better qualitative performance compared with methods trained with all GTs.

Cho 2022; Liu et al. 2022; Yan et al. 2023a,b; Hu et al. 2024; Kong et al. 2024). These methods require paired training samples: LDR images with corresponding HDR ground truths (GTs). The HDR GTs are usually obtained with the help of expensive specially designed imaging systems (Reinhard et al. 2010), or by controlling the scene motion (Kalantari, Ramamoorthi et al. 2017; Chen et al. 2021). This makes it challenging to collect training samples with HDR GTs at scale, which curbs the development of this task. To tackle this issue, recent advances delve into the problem of annotation-efficient HDR image reconstructing: *how to train a HDR reconstructing model with limited HDR GTs?*

The FSHDR (Prabhakar et al. 2021) generates pseudo HDR GTs and inverts them into synthesized LDR images to obtain paired training samples. However, the domain gap between the real and synthesized LDR images may confuse the model, which leads to the artifacts shown in the first column of Fig. 1. The SMAE (Yan et al. 2023c) follows the teacher-student manner and proposes an adaptive pseudo-label selection strategy (APSS) based on the similarity be-

tween the pseudo GTs and the reference images. This strategy ignores the saturation area, which is the key challenge of HDR reconstructing. Besides, the teacher and student in SMAE have a marginal gap. Previous self-/semi-supervised learning methods (Rasmus et al. 2015; Laine and Aila 2016; Tarvainen and Valpola 2017; Sohn et al. 2020a) have demonstrated that the large enough gap between the teacher and student matters for learning new knowledge from unlabeled data. Therefore, the limitations of prior arts can be summarized as: (a) the sub-optimal choice of the teacher model; and (b) the incomplete pseudo GT selecting strategy.

For the limitation (a), we develop a new framework in which the student model is updated by the regular gradient descent, and the teacher model is updated by the exponential moving average (EMA) of student weights. In this manner, the teacher can be seen as a temporal ensembling version of the student. This operation provides a better teacher that can generate high-quality pseudo HDR GTs (Tarvainen and Valpola 2017). In terms of the limitation (b), we propose to mask the pseudo GTs for training according to their reliability rather than hand-crafted rules. Inspired by (He et al. 2019), we add an uncertainty estimation branch to our HDR reconstructing model and append Kullback-Leibler (KL) divergence loss to the original loss. The outputs of the model include not only the HDR images but also uncertainty maps which indicate the reliability of the HDR prediction at pixel level. Based on uncertainty maps, we design a dual threshold policy to mask out the unreliable parts of pseudo GTs at both the patch and pixel levels. This bi-level uncertain area masking policy ensures that reliable and informative pseudo GTs can be learned.

Experimental results illustrate that our framework learns LDR-to-HDR mapping more efficiently with limited HDR GTs: our method can achieve comparable performance with current state-of-the-art fully-supervised methods with only 6.7% HDR GTs, as shown in Fig. 1. The contributions of our work are three-fold:

- We establish a semi-supervised HDR image reconstructing framework firstly taking the predicted uncertainty of pseudo GTs into account.
- We propose a novel bi-level uncertain area masking policy to make sure that trusted pseudo HDR GTs can be learned, which leads to a better performance.
- Our method achieves new state-of-the-art performance under the annotation-efficient setting on public benchmarks.

Related Works

Multi-Frame HDR Image Reconstructing

To reconstruct the ghost-free HDR image from a set of unaligned LDR images with varied exposures, traditional methods align LDR images via homographies (Tomaszewska and Mantiuk 2007) or optical flow (Bogoni 2000; Zimmer, Bruhn, and Weickert 2011); some other works reject (Grosch et al. 2006; Oh et al. 2014; Lee, Li, and Monga 2014) misaligned regions in LDR images before fusing them. However these method leads to unsatisfied results when large motion and saturation occur.

With the rise of the Deep Neural Networks (DNNs), the DNN-based HDR reconstructing methods (Niu et al. 2021; Chung and Cho 2022; Li et al. 2022; Liu et al. 2022; Tel et al. 2023; Cui et al. 2024; Hu et al. 2024; Kong et al. 2024) achieve a significant performance improvement and better visual results. The pioneer works propose to explicitly align LDR images (Liu et al. 2009; Wu et al. 2018) before fusing by CNN network (Kalantari, Ramamoorthi et al. 2017; Wu et al. 2018). However, explicit alignment may fail due to the fact that the misalignment is difficult to be learned. To solve this problem, some works adopt spatial attention to replace the explicit alignment and fuse the attention-based features to reconstruct HDR images (Yan et al. 2019; Liu et al. 2022; Yan et al. 2023c), others propose masking the misaligned regions to mitigate their impact (Kong et al. 2024). Alternatively, diffusion models have been explored to reconstruct HDR images in a generative manner (Yan et al. 2023b; Hu et al. 2024). Nevertheless, these learning-based methods are greatly dependent on the training samples with HDR GTs (Prabhakar et al. 2021). The expensive HDR GTs limit the scale of HDR datasets, which causes a bottleneck.

Therefore, the annotation-efficient methods are introduced to train a HDR reconstructing model with a small amount of HDR GTs (Prabhakar et al. 2021; Yan et al. 2023c). The common strategy of these methods is to make use of both the data with and without HDR GTs. Nevertheless, these methods have neither successfully repaired nor selected high-quality pseudo GTs, nor effectively learned to deghost the HDR results, ultimately resulting in artifact-containing outcomes as shown in Fig. 1. In contrast to previous annotation-efficient methods, we achieve a better performance by masking out the unreliable parts of pseudo GTs at both patch and pixel levels.

Semi-Supervised Learning

Semi-supervised learning (SSL) aims to train models with limited annotations (Qiao et al. 2018; Tarvainen and Valpola 2017; Sohn et al. 2020a; Chen et al. 2022; Wang et al. 2024). Among them, recent methods follow a pseudo-labeling paradigm (Rasmus et al. 2015; Laine and Aila 2016; Sohn et al. 2020a; Chen et al. 2022), which our work also follows. The pseudo-labeling methods train a student model with a mix of human annotations and pseudo annotations generated from a teacher model. These methods significantly enhance various computer vision tasks: detection (Chen et al. 2022; Tang et al. 2021), segmentation (Fan et al. 2022; Wu et al. 2024), and computational photography (Pan et al. 2024a,b).

Nevertheless, when used for new tasks, vanilla pseudo-labeling approaches, such as the mean teacher (Tarvainen and Valpola 2017), can face challenges due to confirmation bias. The confirmation bias means the student model could learn errors from the pseudo annotations (Tarvainen and Valpola 2017). Current approaches aim to mitigate the impact of confirmation bias by carefully selecting pseudo labels (Chen et al. 2022; Liu, Ma, and Kira 2022). Our work addresses the issue of confirmation bias by conducting the pixel-level uncertainty estimation (He et al. 2019) for pseudo HDR GTs. The trusted regions of the pseudo HDR GTs are selected with bi-level thresholds.

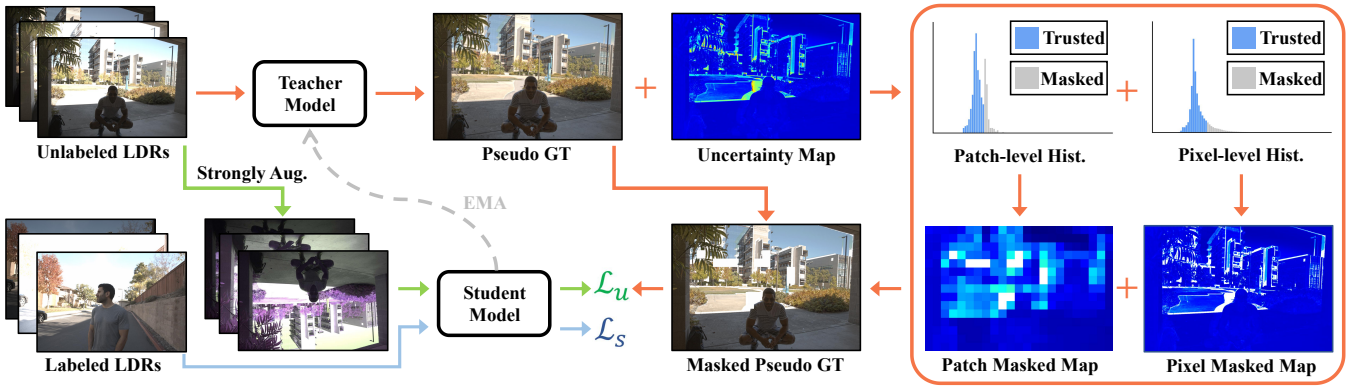


Figure 2. The pipeline of the proposed annotation-efficient HDR image reconstructing framework. The framework follows the teacher-student structure. The teacher branch predicts both the pseudo HDR GTs and the corresponding uncertainty maps. The uncertainty map evaluates the pixel reliability of pseudo GTs. Thus, we can discard the uncertain regions of the pseudo HDR GTs. In this way, the student can hardly be affected by the teacher’s mistakes. We mask out the uncertain regions of the pseudo GTs from both the patch- and pixel-level. The student model learns from the masked pseudo GTs with the unsupervised loss (\mathcal{L}_u). The student also learns from the real GTs with the supervised loss (\mathcal{L}_s). Then, the teacher model is updated by student model via exponential moving average (EMA) (Tarvainen and Valpola 2017).

Methodology

In this section, we firstly introduce the setting of semi-supervised HDR image reconstructing. Secondly, we review the definition of multi-frame HDR image reconstructing. Then, we provide an overview of our teacher-student structure and how to exploit unlabeled data. Finally, we discuss how to select the trusted region in HDR GTs via the proposed bi-level uncertain area masking strategy.

Semi-Supervised Setting

The proposed semi-supervised HDR image reconstructing framework uses paired data $D^l = \{(X^l, y)_i\}_{i=1}^{N^l}$ and unpaired data $D^u = \{(X^u)_i\}_{i=1}^{N^u}$ as training data, where X^l and X^u represents the LDR image burst with and without the corresponding GT HDR image y . N^l and N^u are the total number of data with and without HDR GT, respectively.

Multi-Frame HDR Image Reconstructing

Multi-frame HDR image reconstructing focuses on restoring a ghost-free HDR image from a set of differently exposed LDR images. Each input LDR image set X has three input LDR images (x_1, x_2, x_3) with varied exposures. The middle frame x_2 is treated as the reference image. To project the LDR image set X to the HDR domain, we use gamma correction to generate the gamma-corrected image set \tilde{X} :

$$\tilde{x}_i = \frac{(x_i)^\gamma}{t_i}, \quad i = 1, 2, 3 \quad (1)$$

where t_i indicates the exposure time of x_i , and γ is the gamma correction parameter. Following (Kalantari, Ramamoorthi et al. 2017), we concatenate the original LDR images X and the gamma-corrected images \tilde{X} to obtain the final input $\hat{X} = (X, \tilde{X})$. Both X^l and X^u go through all above operations. The network $\Phi(\cdot)$ is then defined as:

$$r = \Phi(\hat{X}; \theta), \quad (2)$$

where r denotes the reconstructed HDR image, and θ is the network parameter to be optimized. For the network, GFHDR is selected as the teacher and student of our approach, detailed by (Liu et al. 2022).

Model Overview

Fig. 2 shows the proposed framework which follows the teacher-student pseudo-labeling paradigm. This teacher-student framework contains two networks with the same structure but different weights, which are denoted as θ_t and θ_s for the teacher and student, respectively. The teacher branch generates pseudo GTs for the data without HDR GTs. The pseudo GTs are then used as the supervised signal to train the student. Parts of the pseudo GTs contain ghosting or noise, which may misguide the student. Therefore, we design an uncertainty branch to evaluate the quality of the reconstructed HDR images. Then with the uncertainty map, we can mask the unreliable parts of the pseudo GTs. Different from the student being updated by gradients, the teacher is updated by the student weights via the exponential moving average (EMA) (Tarvainen and Valpola 2017):

$$\theta_t^i \leftarrow \alpha \theta_t^{i-1} + (1 - \alpha) \theta_s^i, \quad (3)$$

where α is a hyper-parameter to control the updating speed. θ_t^i and θ_s^i are the parameters of the teacher and student in the i -th iteration, respectively. The student is trained with supervised losses \mathcal{L}_s on images with GTs and unsupervised losses \mathcal{L}_u on images without GTs. As HDR images are normally viewed after tone-mapping $\mathcal{T}(\cdot)$, before computing the losses, both the predicted HDR images and the HDR GTs are tone-mapped by a μ -law function:

$$\mathcal{T}(I) = \frac{\log(1 + \mu I)}{\log(1 + \mu)}, \quad (4)$$

where I can be r or y , μ is a hyper-parameter. Following (Liu et al. 2022), we use \mathcal{L}^r loss and perceptual loss




Figure 3. The architecture of the judge network. The input LDR images are first fed to the attention-based feature extraction network (Liu et al. 2022) to obtain the fused feature F_{att} . Then the three 3×3 convolution layers with a skip addition connection generate the uncertainty map.

\mathcal{L}^v (Johnson, Alahi, and Fei-Fei 2016) as the optimization targets. On the data \widehat{X}^l with real GTs, the supervised losses can be computed as:

$$\mathcal{L}_s^r = \|\mathcal{T}(\Phi(\widehat{X}^l; \theta_s)) - \mathcal{T}(y)\|_1, \quad (5)$$

$$\mathcal{L}_s^v = \sum_j \|\Psi_j(\mathcal{T}(\Phi(\widehat{X}^l; \theta_s))) - \Psi_j(\mathcal{T}(y))\|_1, \quad (6)$$

where $\Psi(\cdot)$ denotes the activation feature maps extracted from a pre-trained VGG-16 network (Simonyan and Zisserman 2014), and j denotes the j -th layer. On the data \widehat{X}^u without GTs, the unsupervised losses can be calculated as:

$$\mathcal{L}_u^r = \|\mathcal{T}(\Phi(\widehat{X}^u; \theta_s)) - \mathcal{T}(y^p)\|_1, \quad (7)$$

$$\mathcal{L}_u^v = \sum_j \|\Psi_j(\mathcal{T}(\Phi(\widehat{X}^u; \theta_s))) - \Psi_j(\mathcal{T}(y^p))\|_1, \quad (8)$$

where $y^p = \Phi(\widehat{X}^u; \theta_t)$ is the pseudo HDR GT. Together with the supervised uncertainty loss \mathcal{L}_s^k and the unsupervised uncertainty loss \mathcal{L}_u^k which will be defined in the following section, the total loss of our framework is:

$$\mathcal{L} = \mathcal{L}_s^r + \lambda_v \mathcal{L}_s^v + \mathcal{L}_s^k + \lambda_u (\mathcal{L}_u^r + \lambda_v \mathcal{L}_u^v + \mathcal{L}_u^k), \quad (9)$$

where λ_u and λ_v are hyper-parameters. The λ_u is used to trade off between the supervised and the unsupervised loss.

Bi-Level Uncertain Area Masking

The challenge of our pseudo-labeling framework lies in how to discard artifacts of the pseudo HDR GTs. We argue to explore trusted pseudo GTs so that the student can learn from the reliable regions of the pseudo GTs. This ensures that the student is unlikely to break down and can perform better. Based on the EMA, the better student leads to the better teacher. So, the updated teacher can produce trusted pseudo GTs on previously uncertain regions. This makes a positive feedback loop. But there is no indicator, *e.g.*, the confidence score in classification task, to evaluate the quality of pseudo labels in this field. Therefore, in this paper, we introduce a judge network to directly provide corresponding evaluation scores along with the predicted HDR pixels. Considering

that the HDR image reconstructing is essentially a dense regression task, we follow the idea of (He et al. 2019) to treat the predictions and the regression targets as Gaussian distributions and Dirac delta functions respectively. Thus, the uncertainty-aware loss for our task can be written as:

$$\mathcal{L}^k = \frac{1}{n_h} \sum_i \left(\frac{(\bar{h}_i - h_i)^2}{2\sigma_i^2} + \frac{1}{2} \log(\sigma_i^2) \right), \quad (10)$$

where the \bar{h} and h are respectively predicted and real/pseudo GT HDR image pixel values, σ_i is the uncertainty score of the i -th pixel, n_h is the total number of pixels. When the value \bar{h} is estimated inaccurately, the network is expected to be able to predict larger σ so that \mathcal{L}^k will be lower.

As shown in Fig. 3, we design a judge network with three-layer 3×3 convolution, one skip connection, and a sigmoid function to predict uncertainty scores. It takes feature F_{att} extracted by the spatial attention module in GFHDR (Liu et al. 2022) as input to obtain the uncertainty map of the predicted HDR image. The uncertainty map explicitly shows the quality of the HDR image by providing the magnitude of the possible deviation of predicted values.

Based on the uncertainty map, we mask out the unreliable regions of the pseudo GTs from the patch and pixel levels to ensure the reliability of the pseudo GTs during utilization. Before training, LDR images and HDR images need to be cropped into patches, we reject the pseudo GT patches that have ghosting or artifacts. In this process, we first calculate the mean value of the uncertainty map at the patch level, then perform global normalization to get the normalized patch-level uncertainty score S_{pa}^i . Furthermore, we obtain a patch-level mask M_{pa}^i according to a pre-set threshold τ_{pa} . The patch-level mask is calculated by:

$$M_{pa}^i = \begin{cases} 1 & S_{pa}^i < \tau_{pa} \\ 0 & S_{pa}^i > \tau_{pa} \end{cases}, \quad (11)$$

where 1 means that the corresponding patch data is used as training data and vice versa.

However, there may still be small areas of ghosting and artifacts in the remaining patches, which requires more refined filtering at the pixel level. Therefore, we first calculate the mean value of each pixel in RGB channel and normalize the values in all pseudo labels to obtain a normalized single-channel uncertainty map. A binary uncertainty mask M_{pi}^i is then generated by a predefined threshold τ_{pi} :

$$M_{pi}^j = \begin{cases} 1 & S_{pi}^j < \tau_{pi} \\ 0 & S_{pi}^j > \tau_{pi} \end{cases}. \quad (12)$$

With the help of this uncertainty mask, we can accurately filter out low-quality pixels with 0 values in M_{pi}^i , which are not used to calculate losses during pseudo label training.

Data Augmentation

Suggested by FixMatch (Sohn et al. 2020b), data augmentation with appropriate differences can lead to a success in SSL by incorporating consistency learning and pseudo-labeling. Therefore, for the ‘‘Strongly Aug.’’ in Fig. 2, we

Dataset	Metric	AHDRNet	ADNet	GFHDR	SCTNet	SAFNet	FSHDR	SMAE	Ours
Kalantari	PSNR- μ	41.05±0.32	40.93±0.38	<u>41.63±0.23</u>	41.58±0.27	37.81±0.44	41.40±0.13	41.61±0.08	44.04±0.07
	PSNR- l	40.61±0.10	40.78±0.15	39.85±0.42	39.93±0.40	37.95±0.65	41.39±0.12	<u>41.54±0.10</u>	41.67±0.06
Hu	PSNR- μ	43.42±0.44	43.79±0.48	42.97±0.17	43.18±0.62	41.47±0.37	43.98±0.27	44.24±0.17	45.10±0.25
	PSNR- l	46.37±0.76	46.88±0.81	45.80±0.93	45.09±0.79	44.39±0.94	47.13±0.13	<u>47.41±0.12</u>	47.93±0.14

Table 1: Results on Kalantari’s (Kalantari, Ramamoorthi et al. 2017) and Hu’s (Hu et al. 2020) dataset under the semi-supervised setting of SMAE (Yan et al. 2023c). The best and the second-best results are respectively highlighted in **bold** and underline.

propose to superimpose RGB random shuffle on top of horizontal flip and 90-degree counterclockwise rotation to generate strongly augmented unlabeled LDRs with pseudo GTs for the unsupervised training of student. A contrast experiment on the composition form of the “Strongly Aug.” is provided in the Supplementary Materials. In addition, to make better use of the data with HDR GTs, we follow GFHDR (Liu et al. 2022) to perform vertical flip and 90-degree clockwise rotation for augmentation.

Details

The proposed semi-supervised HDR deghosting framework is trained with two stages: the warm-up stage and the semi-supervised stage. In the warm-up stage, the student only learns from the labeled data D^l . During the semi-supervised stage, the teacher first initializes weights θ_t via copying the student’s weights θ_s and is then updated by EMA along with the training step of the student. Compared with the teacher θ_t updated after every step, the pseudo HDR GTs are generated after every epoch. Then the trusted regions in pseudo HDR GTs are selected with the guidance of uncertainty maps. Finally, the student can learn from both the annotated data and trusted pseudo-labeled data.

Experiments

Dataset and Settings

Datasets. Following previous methods (Liu et al. 2022; Yan et al. 2023c), we use the training and testing datasets provided by Kalantari (Kalantari, Ramamoorthi et al. 2017) and Hu (Hu et al. 2020). Each of these scenes include three differently exposed input LDR images (exposure values of $\{-2, 0, +2\}$ or $\{-3, 0, +3\}$ for Kalantari’s and $\{-2, 0, +2\}$ for Hu’s) which contain dynamic elements (*e.g.*, camera motion, non-rigid movements) and a GT HDR image aligned with the medium frame captured via static exposure fusion. Moreover, we quantitatively and qualitatively evaluate our method on these two testing sets.

Evaluation Metrics. We mainly include two kinds of objective metrics for evaluation: PSNR and SSIM. In general, we calculate PSNR- l , PSNR- μ , SSIM- l and SSIM- μ to make a quantitative score between the predicted HDR images and corresponding GTs. The $-l$ means the metric is computed directly in the linear domain, While the $-\mu$ means the HDR images are first tonemapped by Eq. 4. The HDR-VDP-2 (Mantiuk et al. 2011) is also calculated, which evaluates the HDR image pairs in both visibility and quality.

Experimental Settings. The proposed framework is implemented by PyTorch. During training, the input images are randomly cropped with size of 64×64 and stride of 32. We use the Adam optimizer for the student, and set the initial learning rate to $2e-4$, β_1 to 0.9, β_2 to 0.999, and ϵ to $1e-8$. We set γ to 2.2, μ to 5000 for the tonemapping, α to 0.999 for the EMA, and λ_u to 1 for the unsupervised loss. The τ_{pa} and τ_{pi} are both set to 0.4 for mask threshold. Sensitivity analyses about the selection of α , τ_{pa} and τ_{pi} are in the Supplementary Materials. The framework is trained for 200 epochs: 30 for warm-up and 170 for semi-supervised training. With $N^l = 5$ and a batch size of 64, training takes about two days on two NVIDIA 3090Ti GPUs. Since the GFHDR is selected as our teacher and student model, our model’s inference time is the same as GFHDR, which can be referred to (Liu et al. 2022) for more details.

Results on Kalantari’s and Hu’s Dataset

Table 1 presents a comparison between our method and several state-of-the-art DNN-based approaches under the semi-supervised setting introduced by SMAE (Yan et al. 2023c). For a fair comparison, all results are reported with means and 95 % margin of variations across 5 runs. As shown in Table 1, our method significantly outperforms the baseline method and attains the SOTA performance on all metrics of two datasets. The proposed method achieves superior performance over the second-best approach, with improvements of 2.41 dB and 0.13 dB in terms of PSNR- μ and PSNR- l on the Kalantari’s dataset, and 0.86 dB and 0.52 dB on the Hu’s dataset, respectively.

More Quantitative Results on Kalantari’s Dataset

The additional experiments on the Kalantari’s dataset in Table 2 are conducted under different semi-supervised settings. Three different semi-supervised settings are chosen depending on the amount of annotated data used: $N^l = 5$ (6.7% GTs), $N^l = 10$ (13.5% GTs), $N^l = 20$ (27% GTs). All the remaining annotated data is treated as unannotated data. The results of supervised SOTA methods trained with total 74 (100% GTs) annotated data are also presented. As shown in Table 2, our method consistently achieves the SOTA performance across all semi-supervised settings. Moreover, our method achieves performance comparable to the fully supervised SOTA methods using only 6.7% labeled data. What’s more, with a total of 27% labeled data, our method achieves the global second-best results in PSNR- l metric. Briefly, all of the above comparison results further confirm the effectiveness of our method.

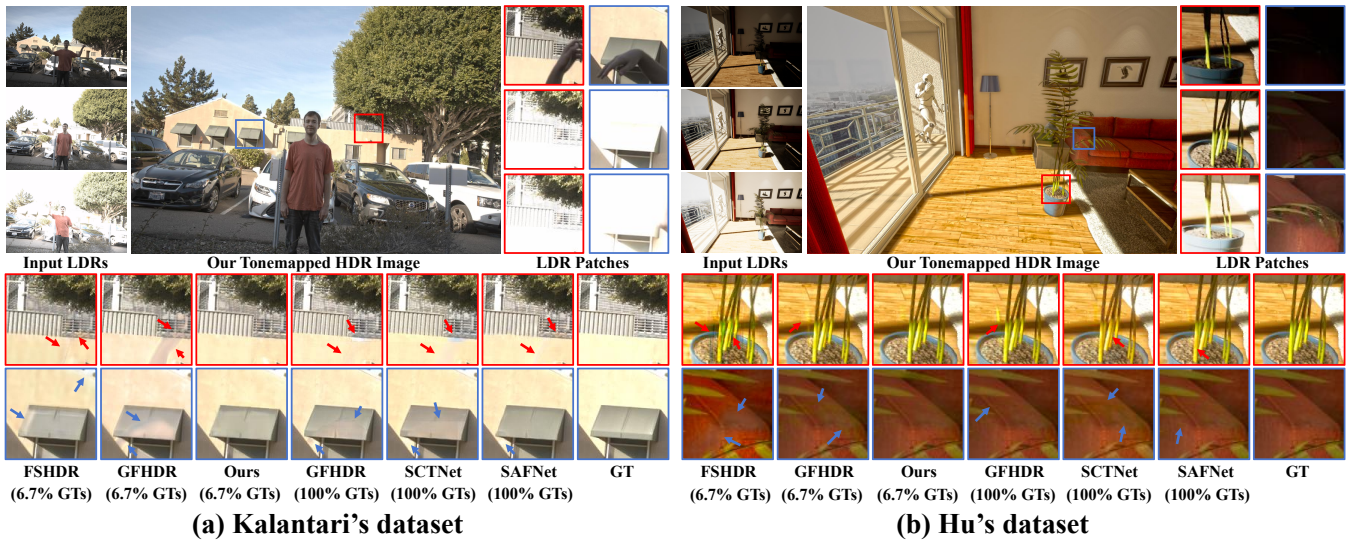


Figure 4. Examples of qualitative results on dataset with GT. Subfigure (a) and (b) present the examples of Kalantari’s (Kalantari, Ramamoorthi et al. 2017) and Hu’s (Hu et al. 2013) datasets, respectively.

Setting	Methods	PSNR- μ	PSNR- l	SSIM- μ	SSIM- l	HV2
100% GTs	Kalantari	42.67	41.23	0.9888	0.9846	65.05
	DeepHDR	41.65	40.88	0.9860	0.9858	64.90
	AHDRNet	43.63	41.14	0.9900	0.9702	64.61
	FSHDR	41.92	41.79	0.9851	0.9876	65.30
	HDR-GAN	43.92	41.57	0.9905	0.9865	65.45
	SMAE	41.97	41.68	0.9895	0.9889	-
	GFHDR	44.32	42.18	0.9916	0.9884	66.03
	HyHDRNet	44.64	42.47	0.9915	0.9894	66.05
	SCTNet	44.49	42.29	0.9924	0.9887	66.65
	DHRNet	44.76	42.59	0.9919	0.9906	66.54
SAFNet	44.66	43.18	0.9932	0.9917	66.93	
6.7% GTs	FSHDR	41.94	40.80	0.9847	0.9860	64.45
	SMAE	41.61	41.54	0.9880	0.9879	-
	GFHDR	42.72	39.80	0.9907	0.9862	63.74
	Ours	44.15	42.55	0.9920	0.9890	64.87
13.5% GTs	FSHDR	42.51	40.97	0.9868	0.9855	64.25
	GFHDR	43.32	40.80	0.9908	0.9867	64.91
	Ours	44.16	42.81	0.9920	0.9899	65.89
27.0% GTs	FSHDR	43.03	41.97	0.9891	0.9863	65.53
	GFHDR	43.14	41.09	0.9911	0.9880	65.17
	Ours	44.21	42.86	0.9921	0.9905	66.10

Table 2: Further semi-supervised and supervised results on Kalantari’s (Kalantari, Ramamoorthi et al. 2017) dataset. The best and the second-best results under different settings are respectively highlighted in **bold** and underline. HV2 represents HDR-VDP-2.

Qualitative Results

As shown in Fig. 4(a) and (b), we conduct qualitative comparisons on Kalantari’s dataset and Hu’s dataset. In fact, due

to the presence of complex factors such as saturation and large motion, achieving color-accurate and ghost-free reconstruction of HDR images is challenging. When the labeled GTs are limited, FSHDR uses reverse optical flow to generate additional pseudo-labeled synthetic data pairs for training. However, the domain gap between the real and synthesized data may misguide the model and lead to unnatural new artifacts and color distortions. Baseline method GFHDR struggles to learn effective alignment from limited GT data, resulting in severe ghosting artifacts. Thanks to the more trusted pseudo GT data and the strong data augmentation-driven consistency learning, our method can learn effective information from unlabeled data, achieving performance close to supervised methods.

Analyses

Ablation studies. As shown in Table 3, we conduct ablation studies on the modules used in our method on Kalantari’s dataset with 6.7% GTs training data to confirm the roles of each module. Among the modules, the GFHDR (Liu et al. 2022) is selected as our baseline (BL) module, and the performance is presented in line i. By comparing the convergence performance with different modules, we can observe:

(a) *The vanilla MT does not work well in our task.* Compared with line i, barely incorporating a vanilla mean teacher (MT) network in line ii without rectifying pseudo GTs does not work. The marginal performance improvement indicates that the model with the only vanilla MT struggles to learn useful information from unlabeled data, which is also known as suffering from confirmation bias. Accordingly, refining the pseudo labels is vital in our semi-supervised framework.

(b) *The uncertainty loss \mathcal{L}_u^k provides uncertainty scores only.* The loss \mathcal{L}_u^k can supervise the learning of the network prediction results and the corresponding scores. The comparison between lines iii and ii shows that the loss function does not affect the results positively or negatively. However,

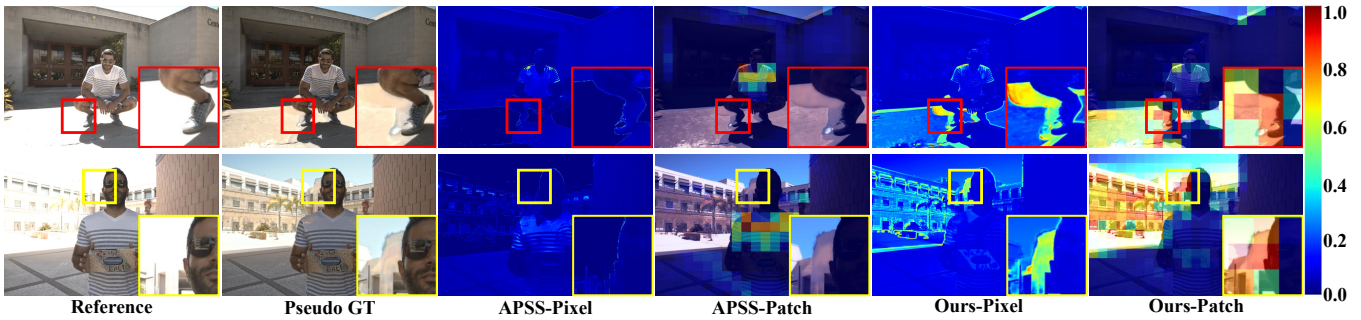


Figure 5. Qualitative comparison of pseudo HDR GT masking strategies. Boxed and zoomed regions highlight notable artifacts in pseudo GTs. Columns 3–6 show APSS uncertainty maps in SMAE and our method at pixel and patch levels. To enhance visual perception, the heat maps are normalized, higher colorbar values indicate more artifacts in pseudo GTs.

No.	BL	MT	\mathcal{L}^k	PaM	PiM	SAug	PSNR- μ	SSIM- μ
i	✓						42.7220	0.9907
ii	✓	✓					42.6509	0.9909
iii	✓	✓	✓				42.8893	0.9905
iv	✓	✓	✓	✓			43.3528	0.9908
v	✓	✓	✓		✓		43.2513	0.9909
vi	✓	✓	✓	✓	✓		43.4787	0.9909
vii	✓	✓	✓	✓	✓	✓	44.0378	0.9915

Table 3: Ablation studies. PaM, PiM and SAug represent patch-level masking, pixel-level masking, and strong augmentations.

as depicted in the last two columns of Fig. 5, the corresponding scores are accurately estimated, which provides a foundation for us to select trusted pseudo GTs.

(c) *The filtration at two levels is better.* With the corresponding pixel-level scores, the pseudo GTs can be filtered to remove regions with artifacts at two levels, namely the pixel level and patch level. In the comparison between the lines iv to vi with line ii, the better convergence performance in line vi indicates that simultaneous masking at both the pixel and patch levels allows for a more precise filtration of pseudo GTs, which can be the better rectification strategy.

(d) *The strong augmentations for pseudo-labeled data training are necessary.* Following (Sohn et al. 2020a), pseudo-labeled data learning in the student needs to be driven by strong augmentations. With the proposed strong augmentations, our method in line vii composed of all modules, achieves a significant improvement over the baseline.

The difference in pseudo GTs scoring strategy. When the teacher model’s capacity is insufficient, the generated pseudo GTs often exhibit ghosting artifacts as highlighted in the second column of the Fig. 5. The SMAE utilizes the similarity between the generated HDR image and the reference frame as the scoring criterion, while ignoring saturated regions. This results in low scores that overlook ghosting artifacts and hinder the learning of well-reconstructed regions, which explains the inferior performance of APSS in line i of Table 4. We propose the judge network and employ the \mathcal{L}^k loss to guide the learning of uncertainty scores, enabling

No.	BL	MT	FSHDR	APSS	Ours	PSNR- μ	SSIM- μ
i	✓	✓		✓		42.7956	0.9910
ii	✓		✓			43.5640	0.9914
iii	✓	✓			✓	44.0378	0.9915

Table 4: Ablation studies. The scoring strategy of the APSS in SMAE and semi-supervised strategy in FSHDR are conducted with the same settings as ours.

the model to assess its own predictions in an uncertainty-aware manner. This facilitates the evaluation of pseudo HDR ground truths and effectively identifies ghosting-affected areas, as illustrated in the fifth and sixth columns of Fig. 5. Consequently, the model’s performance can be effectively improved by leveraging our proposed dual-threshold filtering strategy to select reliable pseudo GTs.

The strategies in semi-supervised learning. In lines ii and iii of Table 4, we present the performance of the BL module trained in FSHDR or ours strategy. FSHDR utilizes reverse optical flow to generate additional synthetic pseudo-labeled data, which, however, disrupts the original distribution of real data. In contrast, our method does not change the original distribution and achieves superior performance.

Conclusion

Paired LDR-HDR data is rare, which limits the development of learning-based HDR image reconstructing. Previous works have attempted to solve this problem from the perspective of semi-supervised learning but struggle to obtain good enough pseudo HDR GTs. To deal with this limitation, we develop a novel teacher-student framework where the student learns from the EMA version of itself, *i.e.*, the teacher model, and propose a bi-level uncertain area masking strategy. In this way, the student model can learn from trusted pseudo labels, which consequently yields a better teacher model. Our proposed method achieves significantly superior performance to the state-of-the-art semi-supervised methods. Moreover, using the data with only 6.7% total available GTs, we achieve comparable performance with state-of-the-art fully supervised methods.

References

- Bogoni, L. 2000. Extending dynamic range of monochrome and color images through fusion. In *Proc. Int. Conf. Pattern Recog.*, volume 3, 7–12. IEEE.
- Chen, B.; Chen, W.; Yang, S.; Xuan, Y.; Song, J.; Xie, D.; Pu, S.; Song, M.; and Zhuang, Y. 2022. Label matching semi-supervised object detection. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 14381–14390.
- Chen, G.; Chen, C.; Guo, S.; Liang, Z.; Wong, K.-Y. K.; and Zhang, L. 2021. Hdr video reconstruction: A coarse-to-fine network and a real-world benchmark dataset. In *Proc. Int. Conf. Comput. Vis.*, 2502–2511.
- Chung, H.; and Cho, N. I. 2022. High dynamic range imaging of dynamic scenes with saturation compensation but without explicit motion compensation. In *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 2951–2961.
- Cui, J.; Jiang, W.; Peng, Z.; Pan, Z.; and Cao, Z. 2024. Exposure Completing for Temporally Consistent Neural High Dynamic Range Video Rendering. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 10027–10035.
- Fan, J.; Gao, B.; Jin, H.; and Jiang, L. 2022. Ucc: Uncertainty guided cross-head co-training for semi-supervised semantic segmentation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 9947–9956.
- Grosch, T.; et al. 2006. Fast and robust high dynamic range image generation with camera and object movement. *Vis. Model. Visual.*, 277284: 3.
- He, Y.; Zhu, C.; Wang, J.; Savvides, M.; and Zhang, X. 2019. Bounding box regression with uncertainty for accurate object detection. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2888–2897.
- Hoefflinger, B. 2007. *High-dynamic-range (HDR) vision*. Springer.
- Hu, J.; Choe, G.; Nadir, Z.; Nabil, O.; Lee, S.-J.; Sheikh, H.; Yoo, Y.; and Polley, M. 2020. Sensor-realistic synthetic data engine for multi-frame high dynamic range photography. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog. Workshop.*, 516–517.
- Hu, J.; Gallo, O.; Pulli, K.; and Sun, X. 2013. HDR deghosting: How to deal with saturation? In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 1163–1170.
- Hu, T.; Yan, Q.; Qi, Y.; and Zhang, Y. 2024. Generating content for hdr deghosting from frequency view. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 25732–25741.
- Johnson, J.; Alahi, A.; and Fei-Fei, L. 2016. Perceptual losses for real-time style transfer and super-resolution. In *Proc. Eur. Conf. on Comput. Vis.*, 694–711. Springer.
- Kalantari, N. K.; Ramamoorthi, R.; et al. 2017. Deep high dynamic range imaging of dynamic scenes. *ACM Trans. Graph.*, 36(4): 144–1.
- Kong, L.; Li, B.; Xiong, Y.; Zhang, H.; Gu, H.; and Chen, J. 2024. Safnet: Selective alignment fusion network for efficient hdr imaging. In *European Conference on Computer Vision*, 256–273. Springer.
- Laine, S.; and Aila, T. 2016. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*.
- Larson, G. W. 1998. Logluv encoding for full-gamut, high-dynamic range images. *J. Graph. Tools*, 3(1): 15–31.
- Lee, C.; Li, Y.; and Monga, V. 2014. Ghost-free high dynamic range imaging via rank minimization. *IEEE Signal Process. Letters*, 21(9): 1045–1049.
- Li, R.; Wang, C.; Wang, J.; Liu, G.; Zhang, H.-Y.; Zeng, B.; and Liu, S. 2022. Uphdr-gan: Generative adversarial network for high dynamic range imaging with unpaired data. *IEEE Trans. Circuits Syst. Video Technol.*, 32(11): 7532–7546.
- Liu, C.; et al. 2009. *Beyond pixels: exploring new representations and applications for motion analysis*. Ph.D. thesis, Massachusetts Institute of Technology.
- Liu, Y.-C.; Ma, C.-Y.; and Kira, Z. 2022. Unbiased teacher v2: Semi-supervised object detection for anchor-free and anchor-based detectors. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 9819–9828.
- Liu, Z.; Lin, W.; Li, X.; Rao, Q.; Jiang, T.; Han, M.; Fan, H.; Sun, J.; and Liu, S. 2021. ADNet: Attention-guided deformable convolutional network for high dynamic range imaging. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 463–470.
- Liu, Z.; Wang, Y.; Zeng, B.; and Liu, S. 2022. Ghost-free High Dynamic Range Imaging with Context-aware Transformer. In *Proc. Eur. Conf. on Comput. Vis.*, 344–360. Springer.
- Mantiuk, R.; Kim, K. J.; Rempel, A. G.; and Heidrich, W. 2011. HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Trans. Graph.*, 30(4): 1–14.
- Niu, Y.; Wu, J.; Liu, W.; Guo, W.; and Lau, R. W. 2021. HDR-GAN: HDR image reconstruction from multi-exposed LDR images with large motions. *IEEE Trans. Image Process.*, 30: 3885–3896.
- Oh, T.-H.; Lee, J.-Y.; Tai, Y.-W.; and Kweon, I. S. 2014. Robust high dynamic range imaging by rank minimization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(6): 1219–1232.
- Pan, Z.; Cui, J.; Wang, K.; Wu, Y.; and Cao, Z. 2024a. Pseudo Label Fusion with Uncertainty Estimation for Semi-Supervised Cropping Box Regression. *IEEE Trans. Multimedia*.
- Pan, Z.; Wang, K.; Wu, Y.; Xiao, L.; Cui, J.; Wang, Z.; and Cao, Z. 2024b. Pseudo-Labeling by Multi-Policy Viewfinder Network for Image Cropping. *arXiv preprint arXiv:2407.01971*.
- Prabhakar, K. R.; Senthil, G.; Agrawal, S.; Babu, R. V.; and Gorthi, R. K. S. S. 2021. Labeled from unlabeled: Exploiting unlabeled data for few-shot deep hdr deghosting. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 4875–4885.
- Qiao, S.; Shen, W.; Zhang, Z.; Wang, B.; and Yuille, A. 2018. Deep co-training for semi-supervised image recognition. In *Proc. Eur. Conf. on Comput. Vis.*, 135–152.

- Rasmus, A.; Berglund, M.; Honkala, M.; Valpola, H.; and Raiko, T. 2015. Semi-supervised learning with ladder networks. *Adv. Neural Inform. Process. Syst.*, 28.
- Reinhard, E.; Heidrich, W.; Debevec, P.; Pattanaik, S.; Ward, G.; and Myszkowski, K. 2010. *High dynamic range imaging: acquisition, display, and image-based lighting*. Morgan Kaufmann.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C. A.; Cubuk, E. D.; Kurakin, A.; and Li, C.-L. 2020a. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Adv. Neural Inform. Process. Syst.*, 33: 596–608.
- Sohn, K.; Zhang, Z.; Li, C.-L.; Zhang, H.; Lee, C.-Y.; and Pfister, T. 2020b. A simple semi-supervised learning framework for object detection. *arXiv preprint arXiv:2005.04757*.
- Tang, Y.; Chen, W.; Luo, Y.; and Zhang, Y. 2021. Humble teachers teach better students for semi-supervised object detection. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 3132–3141.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Adv. Neural Inform. Process. Syst.*, 30.
- Tel, S.; Wu, Z.; Zhang, Y.; Heyrman, B.; Demonceaux, C.; Timofte, R.; and Ginhac, D. 2023. Alignment-free hdr deghosting with semantics consistent transformer. In *Proc. Int. Conf. Comput. Vis.*
- Tomaszewska, A.; and Mantiuk, R. 2007. Image registration for multi-exposure high dynamic range image acquisition. In *Int. Conf. Cent. Eur. Comput. Graph. Visual.*
- Wang, K.; Wu, Y.; Pan, Z.; Li, X.; Xian, K.; Wang, Z.; Cao, Z.; and Lin, G. 2024. Semi-supervised class-agnostic motion prediction with pseudo label regeneration and BEVMix. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 5490–5498.
- Wu, S.; Xu, J.; Tai, Y.-W.; and Tang, C.-K. 2018. Deep high dynamic range imaging with large foreground motions. In *Proc. Eur. Conf. on Comput. Vis.*, 117–132.
- Wu, Y.; Pan, Z.; Wang, K.; Li, X.; Cui, J.; Xiao, L.; Lin, G.; and Cao, Z. 2024. Instance Consistency Regularization for Semi-Supervised 3D Instance Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Yan, Q.; Chen, W.; Zhang, S.; Zhu, Y.; Sun, J.; and Zhang, Y. 2023a. A unified HDR imaging method with pixel and patch level. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 22211–22220.
- Yan, Q.; Gong, D.; Shi, Q.; Hengel, A. v. d.; Shen, C.; Reid, I.; and Zhang, Y. 2019. Attention-guided network for ghost-free high dynamic range imaging. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 1751–1760.
- Yan, Q.; Hu, T.; Sun, Y.; Tang, H.; Zhu, Y.; Dong, W.; Van Gool, L.; and Zhang, Y. 2023b. Toward high-quality HDR deghosting with conditional diffusion models. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(5): 4011–4026.
- Yan, Q.; Zhang, S.; Chen, W.; Tang, H.; Zhu, Y.; Sun, J.; Van Gool, L.; and Zhang, Y. 2023c. SMAE: Few-shot Learning for HDR Deghosting with Saturation-Aware Masked Autoencoders. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 5775–5784.
- Zimmer, H.; Bruhn, A.; and Weickert, J. 2011. Freehand HDR imaging of moving scenes with simultaneous resolution enhancement. In *Comput. Graph. Forum*, volume 30, 405–414.