

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

DRHDR: A Dual branch Residual Network for Multi-Bracket High Dynamic Range Imaging

Juan Marín-Vega^{1,2}

Michael Sloth²

Peter Schneider-Kamp¹

Richard Röttger¹

¹ Department of Mathematics and Computer Science (IMADA), University of Southern Denmark ² Esoft Systems

msl@esoft.com {marin,petersk,roettger}@imada.sdu.dk

Abstract

We introduce DRHDR, a Dual branch Residual Convolutional Neural Network for Multi-Bracket HDR Imaging. To address the challenges of fusing multiple brackets from dynamic scenes, we propose an efficient dual branch network that operates on two different resolutions. The full resolution branch uses a Deformable Convolutional Block to align features and retain high-frequency details. A low resolution branch with a Spatial Attention Block aims to attend wanted areas from the non-reference brackets, and suppress displaced features that could incur on ghosting artifacts. By using a dual branch approach we are able to achieve high quality results while constraining the computational resources required to estimate the HDR results.

1. Introduction

Digital camera sensors have limited capabilities at capturing the rich ranges of luminance values of natural images. In order to produce images that resemble what the human eye can see, common HDR methods utilize multiple frames with different exposure values, to produce a final HDR image that has a higher fidelity with respect to the original scene. Nevertheless, HDR solutions still have to deal with the challenges that arise when merging brackets under different light conditions: Saturated regions and noise. Moreover, camera motion and moving objects can also affect the final estimated image by introducing ghosting artifacts.

Several solutions have been proposed over time to solve these challenges: pixel rejection approaches [4,6,10,11,13, 16,17] that select and minimize the contributions of areas that contain motion objects or misalignment, but also pixel



Figure 1. **Qualitative comparison** from our validation split comparing the AHDR Baseline with our technique. Our solution effectively merges the three input brackets and provide high quality outputs, free of noise or saturated areas. Moreover, effectively achieves higher fidelity on the zoomed areas when compared with the ground truth sample.

registration techniques [1,7,9,22,32] for selecting the areas that would provide the best content for the final solution.

With the recent increase on computational capabilities and the improved collection of data, several deep learning systems trained in a supervised fashion have improved HDR estimation results with a big margin with respect of previous approaches. These deep learning techniques are able to work on a feature space [8, 27], reject and align regions based on reference frames [27,29] and on a non-determined number of frames [2] for producing high fidelity output estimations.

However, despite the impressive quality improvements,

³ Project available at https://github.com/drhdr-user/ drhdr

these techniques still are far from optimal in terms of their computational requirements. On this basis, the NTIRE 2022 HDR Challenge [19] aims to stimulate research on HDR imaging techniques, with an emphasis on efficient solutions. Two tracks are proposed in the competition. Track 1: Fidelity (Low-complexity constrain), and Track 2: Low-Complexity (Fidelity constrain). In Track 1, participants are asked to optimize fidelity scores (PSNR and PSNR- μ) while keeping complexity under 200 GMACs. In Track 2, participants are asked to minimize the complexity of their solutions (GMACs and runtime) while achieving at least the same fidelity scores (PSNR and PSNR- μ) as the baseline method AHDR [29]. In this article, we describe our proposed solution for Track 2.

Inspired by previous works [5, 14, 23, 24, 27] that combine multiple-resolution transformations, either with U-Nets, low resolution branches with operations that are locally smooth when translated to full resolution results; our network aims to leverage on a spatially reduced feature space for alleviating the number of computations that would require operating only on full resolution. Hence, our network is built with two branches that operate on different resolutions. Branch b_0 operates at full resolution while b_1 operates at a fourth of the original resolution. We also incorporate previously successful ideas as the Spatial Attention, Deformable Convolutions, and Dilated Residual Dense Blocks. This combination provides a faster, and less computationally expensive framework, with an increase in quality compared to baseline solutions. We summarize our contributions as follows:

- We propose a dual branch system that works at two different resolution for reducing the computational complexity while still improving on quality metrics.
- We asses the benefits of the components used on each of the branches, a careful explanation of the different training options, followed by quantitative and qualitative results.

2. Related work

For High Dynamic Range Imaging with dynamic scene we have a series of major challenges to solve. Restoring highly saturated regions, noise, misalignment from camera motion on the LDR brackets and dealing with entities in movement. The latest being the most challenging one since foreground objects might occlude regions of interest, but also creating ghosting artifacts. These ghosting artifacts have a similar nature to the ones created by non-aligned brackets, but they can be manifested in a much more severe manner, since objects can have big displacements across the different brackets.

Several approaches have been proposed to tackle these challenging scenarios. Historically, these approaches were

based on either rejection or registration. Pixel rejection approaches aim to detect and minimize the contribution of regions that contain moving objects in order to reduce the probability of ghost areas on the final merged output. Pixels can be rejected after an iterative process of computing weights per pixel and the probability that a pixel is capturing the background [10], predicting and measuring color difference from inputs [4,6], by multi-level thresholding of histograms [13] or by detecting clusters of moving pixels by using binary operations [17]. Lately Lee et al. [11] and Oh et al. [16] proposed using rank minimization for aligning LDR inputs and ghost region detection. Pixel or Patch registration methods aim to estimate HDR outputs by localizing the best regions on the LDR inputs. Bogoni [1] uses Laplacian Pyramids for salience detection, selection and fusion. Kang et al. [9] selects the best pair of LDR based on pixel brightness distribution, and uses hierarchical homography for registration and compensation of pixel motion. The final output incorporates portions from the inputs that are not saturated, and that contain the best details. Jinno and Okuda [7] estimate displacement, occlusion and saturation based on irradiance values of each pixel with a Markov random field model, to produce blur-free results. Zimmer et al. [32] align the input LDR with an energy-based optic flow method that takes into account the varying exposure conditions to create displacement fields with subpixel precition. Sen et al. [22] optimized jointly alignment and reconstruction using a patch-based energy-minimization formulation.

Deep learning approaches have significantly improve results compared to previous techniques. To do so, supervised learning approaches require datasets with input output pairs. Kalantari and Ramamoorthi [8] introduced a convolutional neural network for alignment and merging and collected a dataset of LDR inputs and an HDR output. To enforce the misalignment of the inputs, they replace the original low and high exposure LDR inputs with a pair of low, high LDR inputs that are not aligned with the reference middle bracket. Wu et al. [27] explores the utilization of U-Nets [21] for better exploiting deep representations and perform alignment on the feature space. Yan et al. [29] introduced the use of Spatial Attention Layers, that based on a reference bracket, can detect and suppress regions of the non-reference brackets to produce ghost-free estimations. For NTIRE 2021 Challenge [18], Liu et al. [12] proposed the use of Deformable Convolutions for better alignment of brackets. Previously, the use of Deformable Convolutions had already shown promising improvements on discriminative methods [3, 31] as well as for video frame alignment [25]. Recently, Yan et al. [28] improves upon their previous contributions, by providing a Dual Attention mechanism and a channel Attention Mechanism. Finally, Catley-Chandar *et al.* [2] introduces a system that jointly aligns and merges frames based on uncertainty-drive at-



Figure 2. **Diagram of our method**. Two main branches. b_0 works at full resolution and uses a Deformable Convolutional Block while b_1 works at fourth of the original resolution and uses a spatial attention module. Both branches are supplemented with Dilated Residual Dense Blocks and a Dual Branch Fusion Block.



Figure 3. **Deformable Convolutional Block**: A stack of 2D Convolutions followed by a LeakyReLU are in charge of extracting the guidance features from $[Z_{\gamma}^0, Z_2^0]$. These features are used for generating the offsets and the mask for the deformable convolution. The resultant non-reference features \hat{Z}_{γ}^0 are obtained applying a Deformable Convolution over Z_{γ}^0 together with the mask and the given offsets.

tention maps, and a progressive multi-stage image fusion that can work with an arbitrary number of input brackets. Other deep learning based works have explored unsupervised HDR fusion [20], reinforcement learning for bracket selection [26] and the use of adversarial training for better hallucination of missing content [15].

3. Methodology

3.1. Architecture

The overall architecture comprises two routes, one branch that operates at full resolution, and a second branch that operates at a fourth of the original resolution. We define a fixed number of channels for both branches ch = 42. The full resolution branch adopts a Deformable Convolutional Block [12] while the low resolution branch adopts Spatial Attention [29]. Since both LDR and gamma adjusted contribute to detecting and misalignments [12, 27, 29], we opt for using LDR and gamma adjusted as inputs for both



Figure 4. Spatial Attention Block: Attention maps A_i , i = 1, 3 are obtained from $[Z_{\mathcal{Y}}^1, Z_2^1]$ after 2 2D Convolutions follows by a LeakyReLU and a Sigmoid activation respectively. $\hat{Z}_{\mathcal{Y}}^1$ is finally obtained by a point-wise multiplication on $Z_{\mathcal{Y}}^2$.

branches. Given 3 input LDR brackets I_i , i = 1, 2, 3 the network input is composed as $[I_i, I_i^h]$, i = 1, 2, 3, a channelwise concatenation of the original I and its gamma corrected I^h . Where I_2 is the reference middle bracket, and I_i , i = 1, 3 the non-reference ones.

3.1.1 Full Resolution Branch

The full resolution branch starts with an input layer b_0 composed by a $3 \times 3 \times 42$ 2D Convolution and a Leaky ReLU activation. This layer is responsible of encoding the 6 channels input $[I_i, I_i^h]$ onto 42 channels features Z_i^0 for every bracket i = 1, 2, 3. The input layer b_0 is followed by a Guided Deformable Convolutional Block def and a Dilated Residual Dense Block $[29, 30] DRDB^0$.

Deformable Convolutional Block Similarly to AD-NET [12], the Deformable Convolutional Block def is responsible for aligning images in the feature space. Deformable convolutions improve the ability to model geometric transformations. Given that their spatial support is higher than regular convolutions [31], it can be effective when applied to high resolution, and not so deep, feature representations, accounting for high frequency details and finer alignments. Our def is a light-weight adaptation of the module used in ADNET [12] and EDVR [25]. Instead of a full Pyramid, Cascading and Deformable PCD convolution approach, that works at 3 different resolutions, ours work only on full resolution features. Figure 3 illustrates the Deformable Convolutional Block. A stack of 2D Convolutions followed by a LeakyReLU are in charge of extracting the guidance features from $[Z_{2}^{0}, Z_{2}^{0}]$. These features are used for generating the offsets and the mask for the deformable convolution from which the final \hat{Z}_{2}^{0} is obtained. In particular, 2 def modules are defined, one for each set of non-reference features $l_{0}^{i}, i = 1, 3$.

3.1.2 Low Resolution Branch

The low resolution branch b_1 reduces the spatial dimensionality of Z_i^0 by half on each edge after applying a strided 2D Convolution and LeakyReLU to produce Z_i^1 , with a shape of $B \times 42 \times \frac{H}{2} \times \frac{W}{2}$, where B is the batch size, and H, W the input resolution. Features Z_i^1 are then processed through a Spatial Attention Block and a Dilated Residual Dense Block [29, 30] $DRDB^1$.

Spatial Attention The Spatial Attention Block [29] *att* allows the network to extract features of particular areas of the inputs. It suppresses activations from the features by performing point-wise multiplications. As depicted in Figure 4, the attention maps A_i , i = 1, 3 are obtained from $[Z_{\mathcal{Y}}^1, Z_2^1]$ after 2 2D Convolutions follows by a LeakyReLU and a Sigmoid activation respectively. $\hat{Z}_{\mathcal{Y}}^1$ is finally obtained by

$$\hat{Z_i^1} = A_i \circ Z_i^1, i = 1, 3$$

Two att blocks are defined, for each set of non-reference features $Z_i^1, i = 1, 3$

3.1.3 Dual Branch Fusion

The outputs of $DRDB^0$ and $DRDB^1$ are the features from the branches at different resolutions. dbf is responsible for upscaling the set of low-resolution features from b_1 to match the resolution of b_0 . After this operation, both sets of features are concatenated and fused through a 2D Convolution and a LeakyReLU.

3.1.4 Output Layer

The output layer, uses a global skip dense connection. In particular, the output from dbf and full resolution encoded features Z_2^0 from the reference input $[I_2, I_2^h]$ and apply 2D

Convolution and Leaky ReLU followed by another 2D Convolution and a final ReLU.

4. Experiments

In this section we elaborate on the most important details of the training process and the different experiments that lead to our final solution. We also cover comparisons against baselines. We refer to "Baseline_{in}" for the AHDR implementation trained and evaluated under the same regime as our variants and final solution. We refer to "Baseline" when comparing against the AHDR baseline results from the NTIRE 2022 HDR Challenge [19].

4.1. Training Details

4.1.1 Loss Function

Given that training the network on the tonemapped images is more effective [8, 12, 29], we utilize l_1 loss over the estimated \hat{I}^h and ground truth I^{GT} images. We normalize using using the μ -law tonemapping [12]:

$$mu(x) = \frac{\log(1+\mu x)}{\log(1+\mu)}.$$
 (1)

We utilize $\mu = 5000$ and prior to the tonemapping, we apply tanh normalization using the 99 percentile from the estimated \hat{I}^h .

4.1.2 Data

We split the dataset for Training and Validation. We randomly select 250 samples from the dataset for validation and use the rest for training. The images for training are cropped on 250×250 patches with stride of 250px.

4.1.3 Evaluation Metrics

Evaluation is measured on two metrics. Peask Signal to Noise Ratio PSNR that is computed directly on linear HDR estimations. The second metric is PSNR- μ , which is computed over the estimated images after tonemapping with the μ -law defined in Sec. 4.1.1. The implementation of these metrics are provided by the NTIRE 2022 HDR Challenge [19] administrators.

4.1.4 Training Phases

The network is trained using Adam optimizer with $\beta_1 = 0.9$ and $\beta_1 = 0.999$. The network is trained for 300 epochs in 3 different phases (Figure 6). The training is conducted on a single RTX A6000 with 48GB using a batch size of 28 for approximately 5 days.¹

 $^{^1\}text{By}$ using Nvidia Automatic Mixed Precision AMP during training, memory and time requirements can be reduced by \sim half. However, we



Figure 5. **Qualitative comparison** from our validation split comparing the AHDR Baseline with our technique. Our technique achieves higher accuracy on saturated areas but also produce more clear details.

First training phase The network is initially trained for 140 epochs with lr = 0.0005 and 10 epochs with lr = 0.0001. For speeding up the training process, in this phase we use 50% of the available training data, and randomly replace it at the start of each epoch.

Second training phase The network is trained for 50 epochs using all training data with an starting lr = 0.0001 that decays following Equation 2:

$$decay(x) = x \frac{1-n}{N}.$$
 (2)

where x, n and N are initial learning rate, the epoch and the total of epochs of the current phase.

Last training phase In this phase we train for the remaining 100 epochs using all training data with a decaying learning rate that starts from lr = 0.00005.

4.1.5 Data augmentation

We randomly apply one of the following operations during training: vertical flip, horizontal flip, 90° rotation, or nothing.

4.1.6 Testing

Testing is performed on a RTX Quadro 6000 with 24GB. Our method has a measured runtime of ≈ 0.75 seconds

Method	PSNR ↑	PSNR- $\mu\uparrow$	#weights [↓]	GMACs↓
Variant A	40.95	35.21	1345323	2146.68
Variant B	40.83	35.17	1338087	1526.83
Ours	41.03	35.22	1222035	1769.85
Baseline _{in}	40.61	35.14	1441283	2916.92

Table 1. **Quantitative Results** of the different variants on our validation split. Baseline_{in} refers to the Baseline implementation from AHDR [29] trained and tested under the same regime as the rest. Please refer to Sec. 4.1.7 for more details on the different variants.

at processing images of 1060x1900 pixels. Our method has a peak memory consumption of \approx 14394MiB using float32 precision.

4.1.7 Variants

Variant A We define a Variant A based on the AHDR [29] implementation where we incorporate a low resolution branch with a second Spatial Attention Block and the same number of DRDB blocks (= 6) as in the full resolution branch. In order not to double the number of model weights, we reduce the number of channels on each convolution from ch = 64 to ch = 36. We expect the low resolution branch to compensate the reduction on the number of channels per convolution, but also to further reduce the number of GMACs by performing operations on a spatially reduced feature space.

report results from the model submited to the challenge, that was trained without AMP.

Method	PSNR ↑	PSNR- $\mu\uparrow$
Variant A	38.38	36.89
Variant B	38.35	36.87
Ours	38.50	36.91

Table 2. **Quantitative Results** of the different variants on the validation split from NTIRE 2022 HDR Challenge [19].

Loss	PSNR ↑	PSNR- $\mu\uparrow$
L1	43.85	35.07
$L1_{tanh}$	40.61	35.14

Table 3. **Quantitative Results** of the Base L1 Loss vs $L1_{tanh}$ Loss on the Baseline_{in} model. $L1_{tanh}$ provides better PSNR- μ results. Results computed on our validation split.

Method	PSNR ↑	PSNR- $\mu\uparrow$
Variant B	44.02	35.14

Table 4. **Quantitative Results** of the combining Base L1 Loss vs $L1_{tanh}$ Loss on the Variant B. Results computed on our validation split.

Variant B Based on Variant A, we replace the Spatial Attention Block from the full resolution branch with a Deformable Convolutional Block. We incorporate the Deformable Convolutional Block only on the full resolution branch, to perform a finer alignment and retain high frequency details. On the other hand, the Spatial Attention Block on the low resolution branch, would be responsible for the suppression of non-wanted features from the non-reference brackets. We expect this suppression to still be effective on a spatially reduced feature space. Moreover, since the Deformable Convolutional Block performs a higher number operations (and has a higher number of trainable parameters) compared to the Spatial Attention Block, we compensate it by reducing the number of Dilated Residual Dense Blocks on the full resolution branch, in this case = 3 DRDB on the full resolution branch, = 6 on the low resolution branch.

Final Architecture Table 1 reports results on our validation split. We also include results of Baseline_{in}, the Baseline implementation from AHDR [29] trained and tested under the same regime as Variant A, B and the final architecture *Ours*. It can be seen that Variant A and B, perform above Baseline_{in}. Also, Variant B has similar number of trainable parameters but performs $\approx 25\%$ less operations than Variant A, and $\approx 45\%$ less operations than Baseline_{in}. Still, Variant A achieves slightly better quality performance than Variant B. Motivated by these results, we adopt some



Figure 6. Learning rate training strategy. We perform training on 3 different phases. During the first phase we keep a static training rate of 0.0005 and change it to 0.0001 for the last 10 epochs. Then we train for 2 more phases utilizing a learning rate decay method from Equation 2.

changes to define the final architecture, based on Variant B we increase the number of channels from 36 to 42 and define an equal number of DRDB layers on both branches = 5 (+2 on the full resolution branch, -1 on the low resolution branch), and reduce the number of expanding channels in each DRDB from 36 to 21. This leads to a model with a further reduced footprint, in terms of trainable weights, but a slightly more computationally expensive compared to Variant B. Tab. 1 shows that this final architecture balances out to a smaller but a slightly better architecture overall. We report on Table 2 results on the validation split from the challenge. Notice, these architectures perform above the Baseline implementation from NTIRE 2022 HDR Challenge [19]. It can be seen that the final solution is still more powerful in PSNR and PSNR- μ than Variant A and B. We notice that with this validation split, PSNR- μ results are higher and PSNR results are lower, concluding that our internal validation split reports slightly optimist results for PSNR.

4.2. Loss Function

In order to assess the performance of the loss function, we train a Baseline_{in} implementation with L1 loss, and another using L1_{tanh} where before the tonemapping, we apply tanh normalization using the 99 percentile from the estimated \hat{I}^h . We report the results on Table 3. In this case we see an increase of ≈ 3 dB in PSNR using the L1, but an increase of \approx .7dB on PSNR- μ by using the L1_{tanh}. After observing these results we proceed training Variant B with both losses. We report results on Table 4 on our validation set. We see an improvement on PSNR but a slight degradation on PSNR- μ . We notice that PSNR- μ metric is probably more challenging to increase compared to the regular PSNR: Out of 53 groups contestants, only 14 achieve higher PSNR- μ than the reference implementation. Thus,



Figure 7. Qualitative representation of the behavior from the main components of our system. feat0-h is a single activation channel from Z_3^0 , and dconv-h is the same channel result after applying the Deformable Convlutional Block. Dconv-h represents the effective minimization of details from the over-exposed regions of the table, while enhancing details from the background wall. At the bottom, feat1-1 is a single activation channel from Z_1^0 while att-1 is the same channel after applying the Spatial Attention Block. Att-1 despite being at a fourth of the original resolution, is able to retain the well-exposed details from the table, while representing an accurate reduction of the over-exposed details from the background wall.

LR	PSNR↑	PSNR- $\mu\uparrow$	Epoch
0.0001	40.33	35.04	125
0.0005-0.0001	40.73	35.16	148

Method	Runtime (s)↓	GMACs↓
Ours*	0.972	2534.60
Ours	0.738	1769.85

Table 5. **Quantitative Results** when training our final solution after 150 epochs, following the specifications of Fig. 6 (Only first phase), against a model trained with a learning rate of 0.0001. Training with a slightly higher learning rate regime, the system is able to achieve higher accuracy faster and with room for a higher improvement. Results computed on our validation split.

we decide to proceed only with $L1_{tanh}$ instead of combining both.

4.3. Learning Rate

Finally, we perform a comparison based on the initial learning rate. We compare the performance of our final solution after 150 epochs, following the specifications of Fig. 6 (Only first phase), against a model trained with a learning rate of 0.0001. Table 5 reports the results of this experiment. By training with a slightly higher learning rate our model achieves higher performance. On the other hand, by training with a lower learning rate, the model seems to stabilize on a sub-optimal local minima based on the fact that after epoch 125 the metrics stop improving.

Table 6. **Quantitative Results** when performing inference with an architecture where both branches operate at full resolution (Ours*) on images of 1060x1900 pixels.

4.4. Low Resolution Branch Efficiency

To quantify the efficiency gains by having a branch operating at a fourth of the resolution, we define an architecture (Ours *) where both branches operate at full resolution. Table 6 reports Runtime and GMACs (#weights remain constant) against the final solution. The low resolution branch provides runtime efficiency gains of $\approx 25\%$ and a GMACs reduction of $\approx 30\%$.

4.5. Validation and Final results

For completeness, in Table 7 we include the validation phase results and the final test results of our model and the reference one. Our approach achieves better performance on both metrics, it has less trainable parameters and performs $\approx 40\%$ less operations.

4.6. Qualitative Results

We report qualitative results in Figure 1 and Figure 5 with samples from our validation split. From the different

Method	PSNR ↑	PSNR- $\mu\uparrow$	#weights↓	GMACs↓
Baseline	38.34	36.86	1441283	2916.92
Ours	38.50	36.91	1222035	1769.85
Baseline	37.60	37.02	1441283	2916.92
Ours	38.49	37.11	1222035	1769.85

Table 7. **Quantitative Results**. Validation split (top) and Test split (bottom) results against Baseline [29] implementation from NTIRE 2022 HDR Challenge [19].

images, it can be seen that our method reproduces with high fidelity the ground truth samples. It also equalizes and even outperforms AHDR baseline. Fig. 1 showcases a higher color fidelity while Fig. 5b provides a much clear reconstruction when compared with the ground truth.

Moreover, in Figure 7 we visualize the behavior of the main components of our system. In particular, we illustrate 3 input brackets I_i , i = 1, 2, 3, their respective activations from $[I_i, I_i^h], i = 1, 3$ and the result of Deformable Convolutional Blocks and the Spatial Attention Blocks. feat0-h is a single activation channel from Z_3^0 , and dconv-h is the same channel after applying the Deformable Convlutional Block. On the other hand, feat1-1 is a single activation channel from Z_1^0 while att-1 is the same channel after applying the Spatial Attention Block. Fig. 7 effectively illustrates how the Deformable Convolutional Block suppresses details from the over-exposed regions of the table, while enhancing details from the background wall. For the low exposure bracket, att-1 despite being at a fourth of the original resolution, is able to retain the well-exposed details from the table, while suppressing the over-exposed details from the background wall.

5. Conclusion

In this paper we have proposed DRHDR, an NTIRE 2022 HDR Challenge candidate solution for Track 2: Low-Complexity (Fidelity constrain). We unearth the benefit of exploiting spatially reduced feature representations for alleviating the high computational requirements of full resolution transformations. Despite the lower complexity, we demonstrate the ability of our system to provide accurate, ghost free HDR outputs with superior detail representation and higher efficiency when compared with baseline implementations.

Acknowledgments

This work was supported by the Industrial Ph.D. program's financial support from Innovation Fund Denmark, through the project AIERE (Contract-No: 9065-00099B).

References

- L. Bogoni. Extending dynamic range of monochrome and color images through fusion. In *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, volume 3, pages 7–12 vol.3, Sept. 2000. 1, 2
- [2] Sibi Catley-Chandar, Thomas Tanay, Lucas Vandroux, Aleš Leonardis, Gregory Slabaugh, and Eduardo Pérez-Pellitero. FlexHDR: Modelling Alignment and Exposure Uncertainties for Flexible HDR Imaging. arXiv:2201.02625 [cs, eess], Jan. 2022. 1, 2
- [3] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 764–773. IEEE Computer Society, 2017. 2
- [4] Orazio Gallo, Natasha Gelfandz, Wei-Chao Chen, Marius Tico, and Kari Pulli. Artifact-free High Dynamic Range imaging. In 2009 IEEE International Conference on Computational Photography (ICCP), pages 1–7, Apr. 2009. 1, 2
- [5] Michaël Gharbi, Jiawen Chen, Jonathan T. Barron, Samuel W. Hasinoff, and Frédo Durand. Deep bilateral learning for real-time image enhancement. ACM Trans. Graph., 36(4):118:1–118:12, 2017. 2
- [6] Thorsten Grosch. Fast and robust high dynamic range image generation with camera and object movement. *Vision, Modeling and Visualization, RWTH Aachen*, pages 277–284, 2006. 1, 2
- [7] Takao Jinno and Masahiro Okuda. Motion blur free HDR image acquisition using multiple exposures. In 2008 15th IEEE International Conference on Image Processing, pages 1304–1307, Oct. 2008. 1, 2
- [8] Nima Khademi Kalantari and Ravi Ramamoorthi. Deep high dynamic range imaging of dynamic scenes. ACM Transactions on Graphics, 36(4):1–12, July 2017. 1, 2, 4
- [9] Sing Bing Kang, Matthew Uyttendaele, Simon Winder, and Richard Szeliski. High dynamic range video. ACM Trans. Graph., 22(3):319–325, jul 2003. 1, 2
- [10] Erum Arif Khan, Ahmet Oguz Akyuz, and Erik Reinhard. Ghost Removal in High Dynamic Range Images. In 2006 International Conference on Image Processing, pages 2005– 2008, Oct. 2006. 1, 2
- [11] Chul Lee, Yuelong Li, and Vishal Monga. Ghost-Free High Dynamic Range Imaging via Rank Minimization. *IEEE Signal Processing Letters*, 21(9):1045–1049, Sept. 2014. 1, 2
- [12] Zhen Liu, Wenjie Lin, Xinpeng Li, Qing Rao, Ting Jiang, Mingyan Han, Haoqiang Fan, Jian Sun, and Shuaicheng Liu. ADNet: Attention-guided Deformable Convolutional Network for High Dynamic Range Imaging. In 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pages 463–470, Nashville, TN, USA, June 2021. IEEE. 2, 3, 4
- [13] Tae-Hong Min, Rae-Hong Park, and SoonKeun Chang. Histogram based ghost removal in high dynamic range images. In 2009 IEEE International Conference on Multimedia and Expo, pages 530–533, June 2009. 1, 2

- [14] Sean Moran, Pierre Marza, Steven McDonagh, Sarah Parisot, and Gregory G. Slabaugh. Deeplpf: Deep local parametric filters for image enhancement. *CoRR*, abs/2003.13985, 2020. 2
- [15] Yuzhen Niu, Jianbin Wu, Wenxi Liu, Wenzhong Guo, and Rynson W. H. Lau. HDR-GAN: HDR Image Reconstruction From Multi-Exposed LDR Images With Large Motions. *IEEE Transactions on Image Processing*, 30:3885– 3896, 2021. 3
- [16] Tae-Hyun Oh, Joon-Young Lee, Yu-Wing Tai, and In So Kweon. Robust High Dynamic Range Imaging by Rank Minimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(6):1219–1232, June 2015. 1, 2
- [17] Fabrizio Pece and Jan Kautz. Bitmap movement detection: Hdr for dynamic scenes. In 2010 Conference on Visual Media Production, pages 1–8. IEEE, 2010. 1, 2
- [18] Eduardo Perez-Pellitero, Sibi Catley-Chandar, Ales Leonardis, Radu Timofte, Xian Wang, Yong Li, Tao Wang, Fenglong Song, Zhen Liu, Wenjie Lin, Xinpeng Li, Qing Rao, Ting Jiang, Mingyan Han, Haoqiang Fan, Jian Sun, Shuaicheng Liu, Xiangyu Chen, Yihao Liu, Zhengwen Zhang, Yu Qiao, Chao Dong, Evelyn Yi Lyn Chee, Shanlan Shen, Yubo Duan, Guannan Chen, Mengdi Sun, Yan Gao, Lijie Zhang, Akhil K A, Jiji C V, S M A Sharif, Rizwan Ali Naqvi, Mithun Biswas, Sungjun Kim, Chenjie Xia, Bowen Zhao, Zhangyu Ye, Xiwen Lu, Yanpeng Cao, Jiangxin Yang, Yanlong Cao, Green Rosh K S, Sachin Deepak Lomte, Nikhil Krishnan, and B H Pawan Prasad. NTIRE 2021 Challenge on High Dynamic Range Imaging: Dataset, Methods and Results. In 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pages 691-700, Nashville, TN, USA, June 2021. IEEE. 2
- [19] Eduardo Pérez-Pellitero, Sibi Catley-Chandar, Richard Shaw, Ales Leonardis, Radu Timofte, et al. NTIRE 2022 challenge on high dynamic range imaging: Methods and results. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2022. 2, 4, 6, 8
- [20] K. Ram Prabhakar, V Sai Srikar, and R. Venkatesh Babu. DeepFuse: A Deep Unsupervised Approach for Exposure Fusion with Extreme Exposure Image Pairs. In 2017 IEEE International Conference on Computer Vision (ICCV), pages 4724–4732, Venice, Oct. 2017. IEEE. 3
- [21] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells III, and Alejandro F. Frangi, editors, Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015 - 18th International Conference Munich, Germany, October 5 - 9, 2015, Proceedings, Part III, volume 9351 of Lecture Notes in Computer Science, pages 234–241. Springer, 2015. 2
- [22] Pradeep Sen, Nima Khademi Kalantari, Maziar Yaesoubi, Soheil Darabi, Dan B. Goldman, and Eli Shechtman. Robust patch-based hdr reconstruction of dynamic scenes. ACM Transactions on Graphics, 31(6):1–11, Nov. 2012. 1, 2
- [23] Tamar Rott Shaham, Michaël Gharbi, Richard Zhang, Eli Shechtman, and Tomer Michaeli. Spatially-adaptive pixelwise networks for fast image translation. In *IEEE Confer-*

ence on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021, pages 14882–14891. Computer Vision Foundation / IEEE, 2021. 2

- [24] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep High-Resolution Representation Learning for Visual Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020. 2
- [25] Xintao Wang, Kelvin C.K. Chan, Ke Yu, Chao Dong, and Chen Change Loy. EDVR: Video Restoration With Enhanced Deformable Convolutional Networks. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pages 1954–1963, Long Beach, CA, USA, June 2019. IEEE. 2, 4
- [26] Zhouxia Wang, Jiawei Zhang, Mude Lin, Jiong Wang, Ping Luo, and Jimmy Ren. Learning a Reinforced Agent for Flexible Exposure Bracketing Selection. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 1817–1825, Seattle, WA, USA, June 2020. IEEE. 3
- [27] Shangzhe Wu, Jiarui Xu, Yu-Wing Tai, and Chi-Keung Tang. Deep High Dynamic Range Imaging with Large Foreground Motions. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV* 2018, volume 11206, pages 120–135. Springer International Publishing, Cham, 2018. 1, 2, 3
- [28] Qingsen Yan, Dong Gong, Javen Qinfeng Shi, Anton van den Hengel, Chunhua Shen, Ian Reid, and Yanning Zhang. Dual-Attention-Guided Network for Ghost-Free High Dynamic Range Imaging. *International Journal of Computer Vision*, 130(1):76–94, Jan. 2022. 2
- [29] 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 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 1751–1760, Long Beach, CA, USA, June 2019. IEEE. 1, 2, 3, 4, 5, 6, 8
- [30] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual Dense Network for Image Super-Resolution. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 2472–2481, Salt Lake City, UT, June 2018. IEEE. 3, 4
- [31] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets V2: more deformable, better results. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 9308–9316. Computer Vision Foundation / IEEE, 2019. 2, 3
- [32] Henning Zimmer, Andrés Bruhn, and Joachim Weickert. Freehand HDR Imaging of Moving Scenes with Simultaneous Resolution Enhancement. *Computer Graphics Forum*, 30(2):405–414, Apr. 2011. 1, 2