

# Robust Depth Estimation from Auto Bracketed Images

Sunghoon Im, Hae-Gon Jeon, In So Kweon

Korea Advanced Institute of Science and Technology (KAIST), Republic of Korea

{dlar18927, earboll, iskweon77}@kaist.ac.kr

## Abstract

As demand for advanced photographic applications on hand-held devices grows, these electronics require the capture of high quality depth. However, under low-light conditions, most devices still suffer from low imaging quality and inaccurate depth acquisition. To address the problem, we present a robust depth estimation method from a short burst shot with varied intensity (i.e., Auto Bracketing) or strong noise (i.e., High ISO). We introduce a geometric transformation between flow and depth tailored for burst images, enabling our learning-based multi-view stereo matching to be performed effectively. We then describe our depth estimation pipeline that incorporates the geometric transformation into our residual-flow network. It allows our framework to produce an accurate depth map even with a bracketed image sequence. We demonstrate that our method outperforms state-of-the-art methods for various datasets captured by a smartphone and a DSLR camera. Moreover, we show that the estimated depth is applicable for image quality enhancement and photographic editing.

## 1. Introduction

Many photographers want to capture high-quality images of indoor or night scenes that are insufficiently exposed to light. To do so, they increase exposure time or ISO, but these adjustments can cause other imaging problems, such as motion blur or noise amplification. In an effort to mitigate the physical limitation of camera hardware, several image processing methods have been widely employed, such as single image denoising [6, 4] or edge preserving filtering [29, 13]. However, those approaches often degrade the sharpness of the image or produce cartoonish and surreal results.

The function to take several successive shots with different camera settings called Auto-bracketing (e.g., Exposure, ISO or Flash) or in a very short time called Burst shot has become ubiquitous in most hand-held imaging devices. These photographic techniques for gathering more light have recently attracted interest from the field of com-

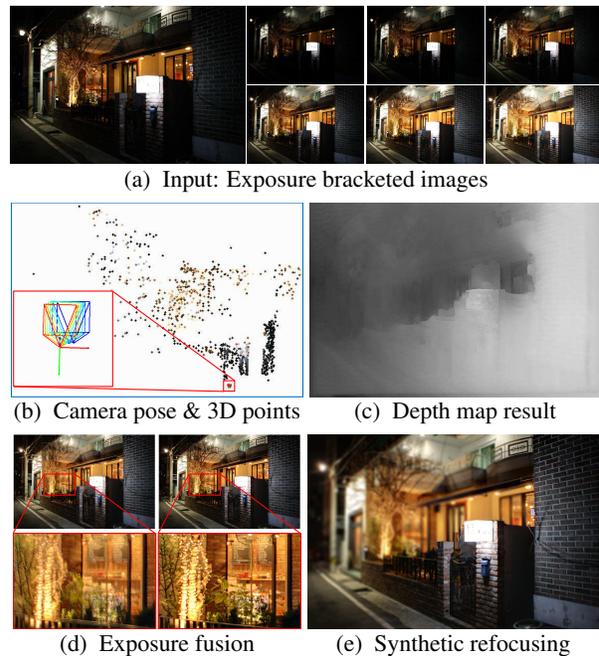


Figure 1: Given exposure bracketed images (a), we estimate camera pose (b) and depth map (c). Our results are applicable to image quality enhancement (d) and depth-aware application (e). We compare exposure fusion results from input images ( $L$ ) and aligned images using our depth ( $R$ ).

putational photography [19, 12]. Assuming that the images are all well-aligned, they are commonly utilized for various image restorations (e.g., Denoising or HDR). However, multiple image alignment is an important issue, since motion inevitably occurs when users press the camera shutter.

In this work, we determined that the inevitable motion, considered as a nuisance in previous burst photography [19, 12], can be used as an important clue to estimate the depth. The estimated depth can be utilized for precise image alignment, which rely highly on discretized homography or optical flow in the conventional methods. Moreover, we show that our depth is useful for various depth-aware applications such as photographic editing or aug-

mented/virtual reality.

Previous studies [32, 15, 8] on the so-called *depth from small motion* (DfSM), have introduced a depth estimation approach based on multiple images with narrow baselines. However, conventional DfSM works have serious limitations, such as (1) noise-sensitive characteristics and (2) high computational complexity, so the estimated depth is not reasonably applicable to hand-held devices as a means of improving image quality. Instead, we propose a learning-based multi-view stereo method combined with the geometric inference.

*Deep neural network* (DNN) has recently been shown to perform well for various computer vision tasks, such as image classification, detection and optical flow estimation. In particular, learning-based optical flow estimation methods [7, 3, 14] outperform conventional optimization-based approaches in accuracy and speed [22]. However, modern geometric interpretations [10] have great advantages in terms of generality and accuracy over the learning-based approaches, *e.g.*, pose estimation [30], and re-localization [16]. To accomplish a robust and fast approach, we complement DNN and modern geometric understandings, and take full advantage of each study.

We first compute a scene geometry including sparse 3D points and camera poses Sec. 2.1, from an input image sequence captured in a burst mode or bracketing mode as shown in Fig. 1a. An output of the scene geometry is then used to obtain a dense depth map by integrating with DNN in Sec. 2.2. Moreover, we show that the estimated depth map can be utilized for precise image alignment in Sec. 2.3. We have carefully evaluated our algorithm using a variety of synthetic and real-world datasets. In the presence of moderate or strong noise and varied intensity in input sequence, our results show considerable improvement over state-of-the-art DfSM methods.

Of course, there are simplified versions of exposure fusion which utilize an image sequence with the same exposure times as the input [12, 11]. Having the same exposures significantly reduces the difficulty in aligning images captured at different times. However, we observe that the burst images can be suffered from many under- or over-exposed pixels when the appropriate exposure time is not determined. The bracketed images are necessary to truly achieve HDR or exposure fusion. We show that our depth can minimize these performance degradations by aligning the images with varying exposures and it is useful for a variety of applications.

## 2. Our Approach

This section describes an effective pipelines for depth and pose estimation method from short burst shots, especially exposure bracketed sequences. First, we introduce robust pose estimation method for intensity variation, which

is slightly modified from the Structure from Small Motion (SfSM) method [15] in Sec. 2.1. Second, we propose a robust depth estimation method tailored for short burst shots even with varied intensity or noise in Sec. 2.2. Lastly, we briefly describe the image alignment method based on our depth and pose information in Sec. 2.3.

### 2.1. Structure-from-Small-Motion (SfSM)

We first extract features from the reference image using Harris corner detection [9], and track the features in a pair of histogram-equalized images using the Kanade-Lucas-Tomasi (KLT) tracker [28]. Before the feature extraction, we perform histogram equalization on all images. Although most commercial cameras have non-linear response functions, this process alleviates the color inconsistency problem in the feature matching step. The equalized images are only used in the feature extraction.

Given the pre-calibrated intrinsic parameters  $\mathbf{K}$ , we estimated the relative camera poses and sparse 3D points by solving the following equation:

$$\underset{\mathbf{r}, \mathbf{t}, \mathbf{X}}{\operatorname{argmin}} \sum_{i=1}^n \sum_{j=1}^m \|\mathbf{u}_{ij} - \langle \mathbf{K}[\mathbf{R}_i | \mathbf{t}_i] \begin{bmatrix} \mathbf{X}_j \\ 1 \end{bmatrix} \rangle\|_2,$$

$$\mathbf{K} = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}, [\mathbf{R}_i | \mathbf{t}_i] = \begin{bmatrix} 1 & -r_i^z & r_i^y & | & t_i^x \\ r_i^z & 1 & -r_i^x & | & t_i^y \\ -r_i^y & r_i^x & 1 & | & t_i^z \end{bmatrix}, \quad (1)$$

where  $\mathbf{r} = [r^x, r^y, r^z]^\top$ ,  $\mathbf{t} = [t^x, t^y, t^z]^\top$  and  $\mathbf{X}_j$  are the rotation, translation components and 3D world coordinates of features.  $\mathbf{u} = [u, v, 1]^\top = \mathbf{K}\mathbf{x}$  and  $\mathbf{x} = [x, y, 1]^\top$  are the image coordinates and normalized image coordinates, respectively.  $n$  and  $m$  are the number of images and features.  $\|\cdot\|_2$  is the L2 norm and  $\langle \cdot \rangle$  is the projection function, that is  $\langle [a, b, c]^\top \rangle = [a/c, b/c]^\top$ .

We initialize all camera components  $\mathbf{r}$ ,  $\mathbf{t}$  to zero and the 3D points  $\mathbf{X}$  by multiplying the normalized image coordinates  $\mathbf{x}$  by a random depth value. We use the Levenberg-Marquardt (LM) optimization [24] to solve Equation (1).

### 2.2. Deep Multi-view Stereo Matching (DMVS)

In this subsection, we describe the detail of our residual-flow network and the derivation of our geometrical transformation that enables to effectively match multiple images. Then, we present our DNN-based multi-view stereo method that incorporates the network and transformation.

**Transformation of optical flow to depth** Rotation alignment reduces the complexity of the transformation between optical flow and depth, which makes our problem more tractable. To disregard the rotational motion, we start by rotating the optical axis of all images to be parallel to that of the reference image. Given the camera intrinsic  $\mathbf{K}$  and

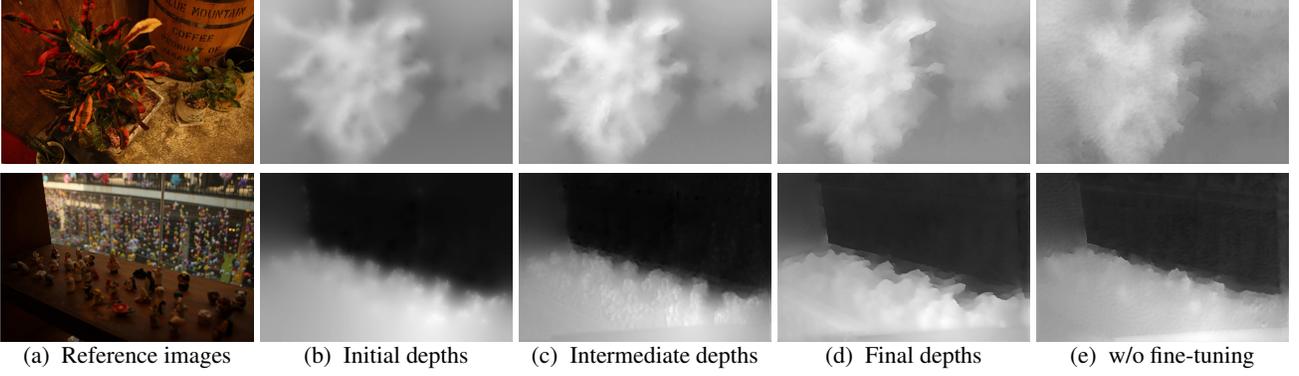


Figure 2: Depth maps according to the number of iterations and fine-tuning. (a) Reference images. (b) The very first initial depths. (c) Intermediate depths from DNN. (d) Our final depths from DNN. (e) Depths from DNN without fine-tuning.

rotation  $\mathbf{R}$  for all images, the synthesized images  $I_i$  can be generated by warping the original images  $\hat{I}_i$ :

$$I_i(\mathbf{u}) = \hat{I}_i(\langle \mathbf{K}\mathbf{R}_i\mathbf{K}^{-1}\mathbf{u} \rangle), i \in \{1, \dots, n\}. \quad (2)$$

We use a bicubic interpolation for this warping process. Occlusion regions are ignored because the baseline of the input images is extremely narrow. All of the images are warped except for the reference image, and the rotationally aligned images are used as the input of DNN. Using the images with pure translation, we can derive the 2D projection of 3D points  $\mathbf{X}_j$  (the multiplication of the normalized image coordinates of the reference frame  $\mathbf{x}_{1j}$  and its depth  $z_j$ ) into the image plane as:

$$\begin{aligned} s\tilde{\mathbf{u}}_{ij} &= \mathbf{K}[\mathbf{I}|\mathbf{t}_i] \begin{bmatrix} \mathbf{x}_{1j}z_j \\ 1 \end{bmatrix} = \begin{bmatrix} u_{1j}z_j + f_x t_i^x + c_x t_i^z \\ v_{1j}z_j + f_y t_i^y + c_y t_i^z \\ z_j + t_i^z \end{bmatrix} \\ &= \begin{bmatrix} z_j & 0 & f_x t_i^x + c_x t_i^z \\ 0 & z_j & f_y t_i^y + c_y t_i^z \\ 0 & 0 & z_j + t_i^z \end{bmatrix} \begin{bmatrix} u_{1j} \\ v_{1j} \\ 1 \end{bmatrix}, \end{aligned} \quad (3)$$

where  $\tilde{\mathbf{u}}$  is the projected image coordinates and  $s$  is the scale factor. Since the  $z$ -axis translation of the image is much smaller than the minimum scene depth ( $t^z \ll z_{min}$ ) [8], we can assume that  $(z_j + t^z)$  is approximately equivalent to  $z_j$  ( $\approx z_j + t^z$ ). The projection matrix in Eq. (3) can be simplified as:

$$\begin{aligned} \begin{bmatrix} \tilde{u}_{ij} \\ \tilde{v}_{ij} \end{bmatrix} &= \begin{bmatrix} 1 & 0 & T_i^x/z_j \\ 0 & 1 & T_i^y/z_j \end{bmatrix} \begin{bmatrix} u_{1j} \\ v_{1j} \\ 1 \end{bmatrix} = \begin{bmatrix} u_{1j} \\ v_{1j} \end{bmatrix} + w_j \begin{bmatrix} T_i^x \\ T_i^y \end{bmatrix}, \\ T_i^x &= f_x t_i^x + c_x t_i^z, \quad T_i^y = f_y t_i^y + c_y t_i^z, \end{aligned} \quad (4)$$

where  $w_j$  is the inverse depth  $1/z_j$ . Based on Eq. (4), the transformation vector  $\mathbf{T}_i$  can convert the inverse depth vec-

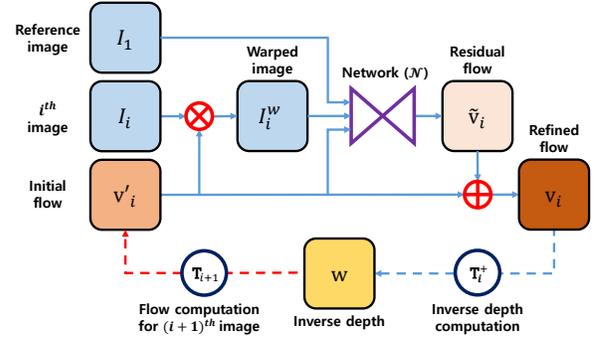


Figure 3: Overview of DMVS. The solid line shows the optical flow refinement process. The blue dashed line shows the conversion of optical flow into inverse depth. The red dashed line shows the inverse depth, which is converted into an optical flow and used for the initial flow of the next frame.

tor  $\mathbf{w}$  into the flow field  $\mathbf{v}_i$  from the reference image to  $i$ -th target image as follows:

$$\mathbf{v}_i = \mathbf{T}_i \mathbf{w}, \quad \text{where } \mathbf{T}_i = \begin{bmatrix} T_i^x \\ T_i^y \end{bmatrix}. \quad (5)$$

**Depth estimation using residual flow network** The basic idea of our depth estimation scheme is to iteratively refine the inverse depth  $\mathbf{w}$  using the optical flow estimated by the DNN as shown in Fig. 3. The network  $\mathcal{N}$  computes the residual flow  $\tilde{v}_i$  with the 8-channel input: the reference image  $I_1$ , the warped image  $I_i^w$  and the initial optical flow  $\mathbf{v}'_i$ . The initial flow is obtained by propagating the sparse 3D points in Sec. 2.1 using the closed form solution [15], which then is transformed into a flow field. We obtain the warped image using the bilinear sampler  $S(I_i, \mathbf{v}'_i)$ . After the residual flow is estimated, the initial flow and the residual flow are added to obtain the refined flow  $\mathbf{v}_i$ . We convert



Figure 4: Averaged image of input exposure bracketed images, our depths and example of photographic applications (denoising, exposure fusion) using aligned images.

the refined flow to the flow of the next frame using the transformation vectors by utilizing them as an initial flow:

$$\mathbf{v}'_i = \mathbf{T}_i \mathbf{T}_{i-1}^+ \mathbf{v}_{i-1}, \quad (6)$$

where  $\mathbf{T}^+$  is the pseudo inverse of the vector  $\mathbf{T}$ . We estimate the final depth  $\mathbf{z}$  by transforming the optical flow of the last image  $\mathbf{T}_n^+ \mathbf{v}_n$  into the inverse depth  $\mathbf{w}$  and dividing it by one.

Fig. 2 shows the effectiveness of the refinement process. The initial depth maps in Fig. 2b show inaccurate depth discontinuities, which is not suitable for the precise image alignment and other depth-aware photographic applications. On the other hand, the intermediate and final depth in Fig. 2c and Fig. 2d shows that our DNN produces more detailed and artifact-free depth results.

**Training and network architecture** Our network consists of two convolution and three deconvolution layers with the fixed kernel size ( $7 \times 7$ ) and stride (1) as described in Table 1. All layers with the exception of the last layer are followed by a Rectified Linear Unit (ReLU). Taking a coarse-to-fine strategy similar to the optical flow estimation, we train the network to learn residual flow  $\tilde{\mathbf{v}}$ , instead of directly estimating the depth or optical flow. We stack the reference image, the warped pair image and the initial optical flow to form an 8-channel input for our network. We set the target residual flows  $\tilde{\mathbf{v}}_i^{gt}$  as the difference between the target flow  $\mathbf{v}_i^{gt}$  and the optical flow  $\mathbf{v}'_i$  obtained from the trained network at the 5<sup>th</sup> pyramid level [25]:

$$\tilde{\mathbf{v}}_i^{gt} = \mathbf{v}_i^{gt} - \mathbf{v}'_i. \quad (7)$$

In the training step, we minimize the average endpoint error (EPE), which is the standard error measure for optical flow estimation. This is the Euclidean distance between the residual flow  $\tilde{\mathbf{v}}_i$  and the target residual flows  $\tilde{\mathbf{v}}_i^{gt}$ .

The optimization is carried out using ADAM [17] with the recommended parameters  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ .

Table 1: Specification of our architecture

Name	Kernel	Str.	Ch I/O	Input
conv1	$7 \times 7$	1	8/32	Images/Flow
conv2	$7 \times 7$	1	32/64	conv1
deconv2	$7 \times 7$	1	64/32	conv2
deconv1	$7 \times 7$	1	32/16	deconv2
deconv0	$7 \times 7$	1	16/2	deconv1

The initial learning rate is  $\lambda = 1e-4$ , then decreased to  $1e-5$  after 60 epochs. We use the Flying Chairs dataset [7] with a resolution of  $384 \times 512$  at training time. The training is performed with a customized version of Torch7 [5] on a Nvidia 1080 GPU, which usually takes 24 hours.

We chose to perform various types of data augmentation during training. We perform spatial (rotation, scaling) and chromatic transformations (color, brightness, contrast, Gaussian noise). We augment input patches with random rotations within  $[-17^\circ, 17^\circ]$  and scaling within  $[1, 2]$ . The noise level is uniformly sampled from  $\mathcal{N}(0, 0.1)$ . We also apply color jitter with additive brightness, contrast and saturation sampled from a Gaussian,  $\mathcal{N}(0, 0.4)$ . At the end, we normalize the intensity of the images using a mean and standard deviation computed from a large corpus of ImageNet [26].

The trained network produces accurate residual flow on images captured with constant camera settings, but it causes some artifacts with a different setting (*e.g.*, exposure, ISO) as shown in Fig. 2e. To alleviate this problem, we fine-tune the network using the different color jitter value  $\mathcal{N}(0, 0.4)$  in the reference  $I_1$  and the target images  $I_i$ . The fine-tuning step generates a synthetic image pair with the different camera settings (*e.g.*, exposure, ISO). We also apply the other data augmentation and intensity normalization in this fine-tuning step using the learning rate  $\lambda = 1e-5$ . Fig. 2d shows the performance improvement in the network fine-tuning.

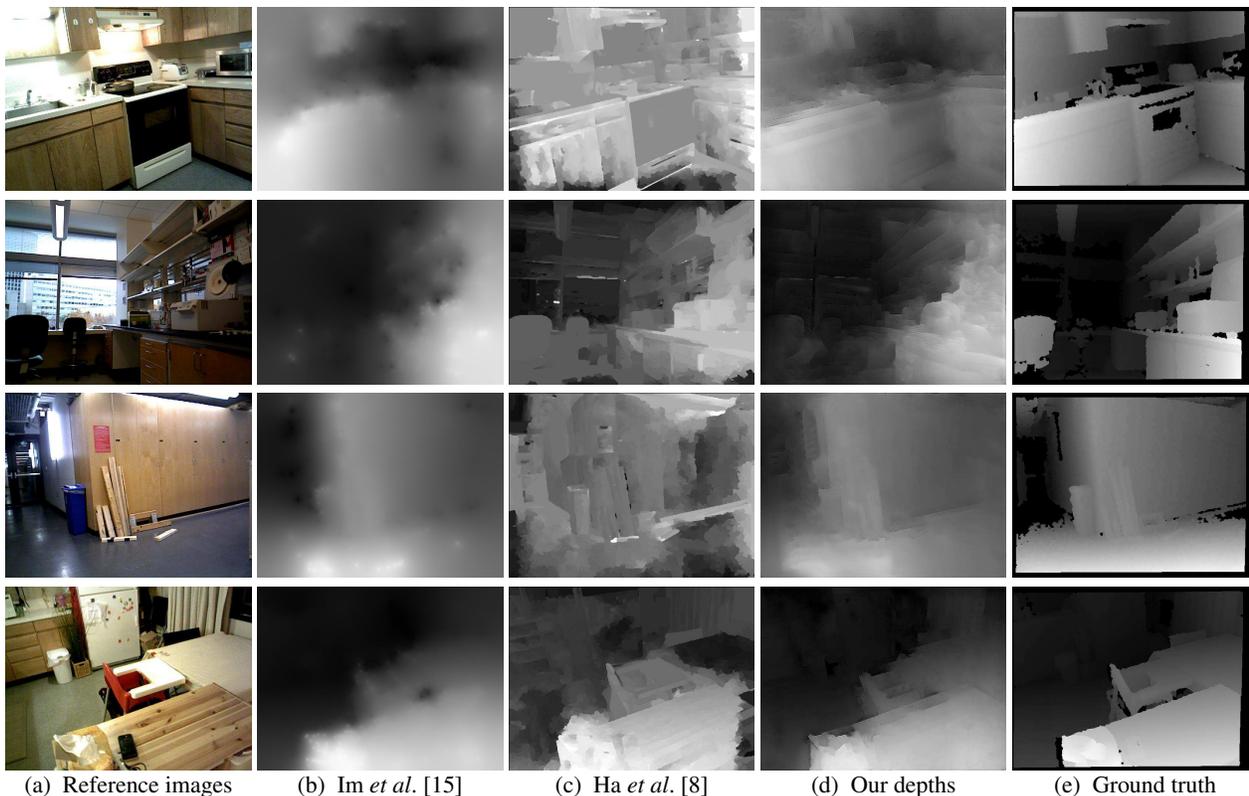


Figure 5: Depth map results using SUN3D datasets [31]. (a) Reference images. (b) Depth maps from propagation [15]. (c) Depth maps from plane sweeping [8]. (d) Our depth maps. (e) Kinect depth maps. *mit\_w85k1*, *mit\_lab\_koch*, *mit\_lab\_16* and *mit\_w85h* (top-to-bottom).

### 2.3. Image alignment

Using the camera geometry  $\mathbf{K}$ ,  $[\mathbf{R}|\mathbf{t}]$  and scene geometry  $\mathbf{z}$  estimated in Sec. 2.2, we can simply align all images. The aligned images  $\tilde{I}_i$  where the original image  $\hat{I}_i$  appears to have been taken at the reference view point are formulated as:

$$\tilde{I}_i(\mathbf{u}) = \hat{I}_i(\langle \mathbf{K}[\mathbf{R}_i|\mathbf{t}_i] \begin{bmatrix} \mathbf{x}_1\mathbf{z} \\ 1 \end{bmatrix} \rangle), i \in \{1, \dots, n\}. \quad (8)$$

We use a bicubic interpolation in this warping process. The aligned images can be used for image quality enhancement applications such as noise reduction and exposure fusion as shown in Fig. 4. Using the estimated depth in Fig. 4b, we warp all non-reference images in Fig. 4a into the reference view point. After aligning the images, we use simple weighted averaging method [19] for denoising in Fig. 4d and exposure fusion algorithm [23] in Fig. 4e. The results show that our estimated depth and pose can precisely align the input images, which is applicable for image quality enhancement.

### 3. Experimental Results

In this section, we demonstrate the effectiveness and robustness of the proposed method using various experiments. First of all, we compared our depth map results to those obtained from the state-of-the-art DfSM methods [15, 8]. In quantitative evaluation, we generated synthetic noisy images from the public RGB-D datasets [31] and utilized them as the input. We then demonstrate that our method produces accurate depth with varying exposure image sequences captured by the bracketing mode. Finally, we investigated the applicability of the depth results for depth-aware photographic application, as well as image quality enhancement.

All steps were implemented in MATLAB<sup>TM</sup>, except for the DNN part, which was implemented by Lua. We set the random depth value to 100, with the constants  $c_1$ ,  $c_2$  and  $\sigma_c$  as 1, 10 and 0.2, respectively. On average, for an image sequence of 28 frames with  $640 \times 480$  resolution, our method took 4s in total for pose and depth estimation on an Intel i7 3.40GHz CPU and 16GB RAM. The SfSM (including feature extraction and bundle adjustment) and the DMVS (including depth propagation and geometric transformation) required 2.5s and 1.5s, respectively.

### 3.1. Synthetic datasets

**Quantitative evaluation of our DMVS** We quantitatively compared our developed approach with the state-of-the-art DfSM methods [15, 8], using public RGB-D datasets [31]. For the datasets, Microsoft Kinect was used to capture the sequential images and the corresponding depth maps. We used 28 consecutive frames for the comparison (previous works requires about 30 frames as input). Since the datasets are taken moving slowly at 30fps, the baseline of the input sequence is narrow enough for quantitative evaluation of DfSM. To simulate realistic camera noise, we applied a signal-dependent Gaussian noise [27] with a standard deviation  $\sigma$  of 0.02. The noise level was determined by averaging the computed noise levels [18] in low-light conditions using a Nexus6.

Fig. 5 shows the depth maps from [15], [8], our method and Microsoft Kinect using the synthetic noisy sequence. As shown in Fig. 5b, [15] fails to show promising results due to inaccurate initial matching cost and the dense depth reconstruction. Work in [8] shows relatively accurate depth discontinuity, but also yields inaccurate depths as shown in Fig. 5c. This is because the plane sweeping algorithm using color similarity as a matching cost is not suitable for images with varying exposures, which produces an unreliable depth map. On the other hand, our DNN-based approach in Fig. 5d has the ability to handle the intensity changes, and to infer an accurate dense depth map, unlike [15, 8].

For a more detailed analysis, we measure a bad pixel rate and *Root-mean-square-error* (RMSE) with varying noise levels ( $\sigma = 0, 0.02, 0.05$ ). *Bad pixel rate* denotes the percentage of pixels that have a distance error of less than 10% of the maximum depth value in the scene. We excluded the unmeasured depth regions due to the hardware limitations of Microsoft Kinect (dark areas in the Kinect depth maps in Fig. 5e) in the error measurement. The results of test across datasets in Fig. 6 shows our method has less RMSE and bad pixel rate than both the state-of-the-art methods for all noise levels. We can see that the conventional methods give acceptable results when noise is not issue, but as noise increases, these measures degrade rapidly. Compared to the competing methods, our method achieved the best results regardless of noise levels, with the least degradation of performance.

### 3.2. Real-world datasets

**Qualitative evaluation of our DMVS** We designed a real-world experiment to verify that the proposed method could be applied to actual exposure bracketed images. First of all, we performed a qualitative comparison of DfSMs [15, 8] using exposure bracketing sequences. We took 28 frames with 7 exposure levels for one second in a commercial DSLR camera (Canon 1D Mark III). Since the state-of-the-art methods have not considered intensity changes,

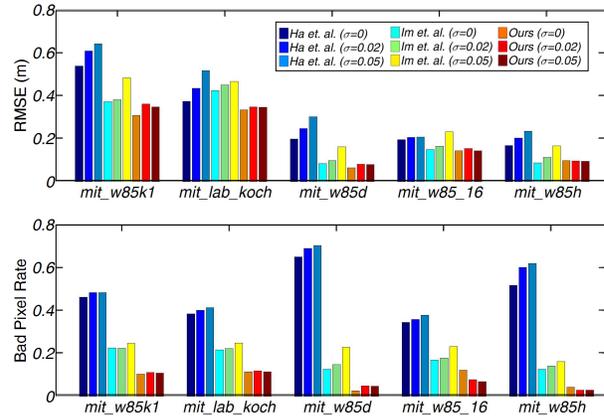


Figure 6: Quantitative evaluation results with state-of-the-art DfSM methods.

we equalized the histogram of all images to adjust image intensity and used them as an input of the methods. Raw images were used for our input.

Fig. 7 shows the results of real-world datasets captured at night. All the comparative methods produce reasonable results; however, we found that our method achieves more reliable results. The propagation method [15] results in over-smoothing effect in Fig. 7b, and the plane sweep method [8] exhibits the speckle artifacts and quantization errors in Fig. 7c. Although brightness-adjusted images were used for the competing methods, over or under-saturation regions might exist, which causes severe artifacts. Despite intensity changes on images, our results in Fig. 7d show an immunity towards the changes, similar to the result of the synthetic datasets in Sec. 3.1.

We also found that our accurate depth can be additionally useful for exposure fusion and depth-aware photographic editing applications, such as digital refocusing and image stylization in Fig. 8. Exposure fusion assembles the multi-exposure sequence into a high quality image using a weighted blending of the input images [23]. To obtain a desirable result, the set of images should be well-aligned. The final results in Fig. 8b demonstrated that our depth can accurately align the set of images. Digital refocusing, which shifts the in-focus region after taking a photo [1, 2], is one of the most popular depth-aware applications. For a realistic refocused image, accurate depth information is necessary. We added synthetic blurs to the images and produced a shallow depth of field image using our depth in Fig. 8c (top). Another interesting application is image stylization, which photographically changes an image at a certain depth range in Fig. 8c (bottom). These results demonstrate that our depth is enough to be utilized on real-world images for various photographic applications.

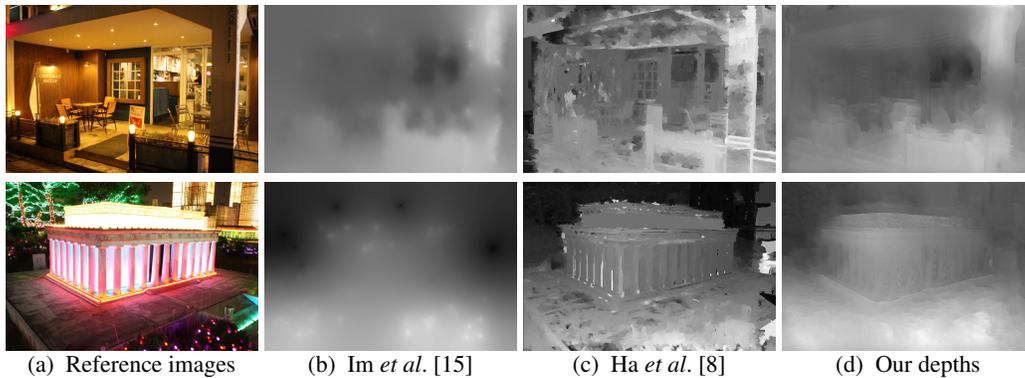


Figure 7: Comparison of depth estimation results with state-of-the-art methods captured by Canon 1D Mark III. (a) Reference images. (b) Depth maps from propagation [15]. (c) Depth maps from plane sweeping [8]. (d) Our final depth maps.

**Comparison to state-of-the-art burst photography** Finally, we compared the proposed method to state-of-the-art burst image photographic approaches; Burst Image Denoising [19] and HDR+ [12]. *Microsoft selfie app* and *Google camera app* pioneered the use of Burst Image Denoising and HDR+ on *iOS8* and *Android*, respectively. We took the image sequences from each phone to use them as the input of our algorithm, and compared them with Burst Image Denoising and HDR+. We obtained independent results from an iPhone5S and Nexus6, as shown in Fig. 9.

Burst Image Denoising [19] aligns the input image sequences using the local homography, then merges them with the weighted average. The denoising results in Fig. 9a shows that Burst Image Denoising outputs blurred results, while our method preserves image boundaries and fine detail in Fig. 9b. (Note that the blurred frame is not our selection, but is the result of image alignment [19].) The local homography might sometimes fail to handle the user’s inevitable motion during burst mode.

The HDR+ [12] generates synthetic exposures by applying gain and gamma corrections to multiple images using a constant exposure, then fuses the synthetic images as if they had been captured using bracketing. Although HDR+ shows promising results in well exposed areas, the constant exposure does not help to recover some badly exposed areas due to lack of light, as shown in Fig. 9c. On the other hand, exposure fusion with real bracketing can cover all of the areas of the input image, as shown in Fig. 9d. Our depth estimation method enables the fusing of bracketed images with exposures that are not aligned, and results in brighter images than the original images.

## 4. Discussion

We have presented a robust narrow-baseline multi-view stereo matching method for noise or intensity changes. We determined an important clue that the baseline of the in-

evitable motion can be used for depth estimation, and the depth enables accurate image alignment leading to image quality enhancement. Both depth and image enhancement results were compared against state-of-the-art methods with a variety of datasets, and demonstrated considerable improvement over existing methods.

The main advantage of our method is its fast computational time and small size network, which are important features for implementation in a mobile platform. Compared to state-of-the-art DfSM methods [15, 8], which take about a few minute, our method takes only a few second. Our DNN plays a key role in reducing computational complexity in dense matching which is the most time-consuming part of conventional DfSM. In addition, our network is much lighter than the DNN-based fast depth or optical flow estimation methods [33, 21, 20, 7] (Flownet: 32M vs Ours: 240K). This significant reduction without performance degradation is achieved by training the residual flow, and iteratively updating optical flow. We expect that the proposed framework will become popular as a mobile phone application.

On the other hand, there are still rooms for improvements: 1) when there is large camera rotation, inaccurate camera poses might be obtained, which can cause an error in our DMVS; 2) our method requires the pre-calibrated intrinsic parameters to estimate the camera poses; 3) the performance of our method is not guaranteed for datasets with fast moving objects, since the scene flow contains additional flow on the object; 4) various fields, such as AR/VR, require metric scale depth, but the estimated depth is not represented in the metric scale.

As future works, we have plan to address these issues. In particular, an idea of the uncalibrated DfSM in [8] is expected to provide a solution to the calibration issue. The scale problem can also be addressed if we directly measure the camera motion during taking photos by introducing additional hardware such as inertial sensors.



Figure 8: Depth-aware photographic editing applications to Synthetic refocusing (top), Image stylization (bottom) and our depths captured by Canon 1D Mark III

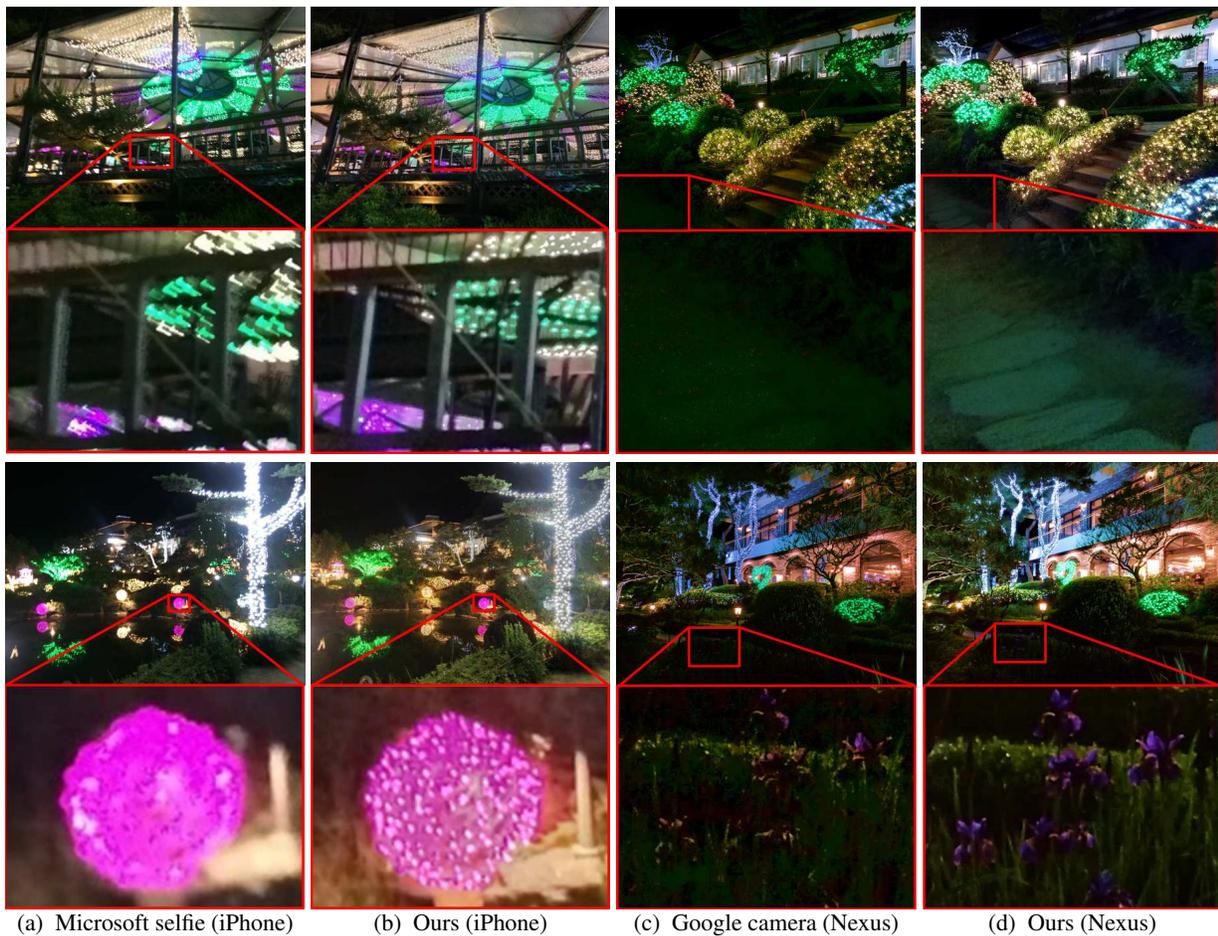


Figure 9: Qualitative comparison with the state-of-the-art methods [19, 12]. (a) Burst images Denoising results from Microsoft selfie app [19]. (b) Our noise-free exposure fusion results. (c) HDR+ results from Google camera app [12]. (d) Our noise-free exposure fusion results. (a), (b) are captured by an iPhone 5S and (c), (d) are captured by a Nexus 6.

**Acknowledgement** This work was supported by the Technology Innovation Program (No. 2017-10069072) funded By the Ministry of Trade, Industry & Energy (MOTIE, Korea). Sunghoon Im was partially supported by Global Ph.D. Fellowship Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (NRF-2016907531).

## References

- [1] HTC One (m8). <http://www.htc.com/us/smartphones/htc-one-m8/>.
- [2] Venue 8 7000 series. <http://www.dell.com/en-us/shop/productdetails/dell-venue-8-7840-tablet/>.
- [3] M. Bai, W. Luo, K. Kundu, and R. Urtasun. Exploiting semantic information and deep matching for optical flow. In *Proc. of European Conf. on Computer Vision (ECCV)*. Springer, 2016.
- [4] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2005.
- [5] R. Collobert, K. Kavukcuoglu, and C. Farabet. Torch7: A matlab-like environment for machine learning. In *BigLearn, NIPS Workshop*, number EPFL-CONF-192376, 2011.
- [6] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007.
- [7] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *Proc. of Int’l Conf. on Computer Vision (ICCV)*, 2015.
- [8] H. Ha, S. Im, J. Park, H.-G. Jeon, and I. So Kweon. High-quality depth from uncalibrated small motion clip. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [9] C. Harris and M. Stephens. A combined corner and edge detector. In *Alvey vision conference*, volume 15, page 50. Manchester, UK, 1988.
- [10] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [11] S. W. Hasinoff, F. Durand, and W. T. Freeman. Noise-optimal capture for high dynamic range photography. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2010.
- [12] S. W. Hasinoff, D. Sharlet, R. Geiss, A. Adams, J. T. Barron, F. Kainz, J. Chen, and M. Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Trans. on Graph.*, 35(6):192, 2016.
- [13] K. He, J. Sun, and X. Tang. Guided image filtering. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(6):1397–1409, 2013.
- [14] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. *arXiv preprint arXiv:1612.01925*, 2016.
- [15] S. Im, H. Ha, G. Choe, H.-G. Jeon, K. Joo, and I. So Kweon. High quality structure from small motion for rolling shutter cameras. In *Proc. of Int’l Conf. on Computer Vision (ICCV)*, 2015.
- [16] A. Kendall, M. Grimes, and R. Cipolla. PoseNet: A convolutional network for real-time 6-dof camera relocalization. In *Proc. of Int’l Conf. on Computer Vision (ICCV)*, pages 2938–2946, 2015.
- [17] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [18] X. Liu, M. Tanaka, and M. Okutomi. Single-image noise level estimation for blind denoising. *IEEE transactions on image processing*, 22(12):5226–5237, 2013.
- [19] Z. Liu, L. Yuan, X. Tang, M. Uyttendaele, and J. Sun. Fast burst images denoising. *ACM Trans. on Graph.*, 33(6):232, 2014.
- [20] W. Luo, A. G. Schwing, and R. Urtasun. Efficient deep learning for stereo matching. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [21] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [22] M. Menze and A. Geiger. Object scene flow for autonomous vehicles. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 3061–3070, 2015.
- [23] T. Mertens, J. Kautz, and F. Van Reeth. Exposure fusion: A simple and practical alternative to high dynamic range photography. In *Computer Graphics Forum*, volume 28, pages 161–171. Wiley Online Library, 2009.
- [24] J. J. Moré. *The Levenberg-Marquardt algorithm: implementation and theory*. Springer, 1978.
- [25] A. Ranjan and M. J. Black. Optical flow estimation using a spatial pyramid network. *CoRR*, abs/1611.00850, 2016.
- [26] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *Int’l Journal of Computer Vision*, 115(3):211–252, 2015.
- [27] Y. Y. Schechner, S. K. Nayar, and P. N. Belhumeur. Multiplexing for optimal lighting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(8):1339–1354, 2007.
- [28] C. Tomasi and T. Kanade. Detection and tracking of point features. 1991.
- [29] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proc. of Int’l Conf. on Computer Vision (ICCV)*. IEEE, 1998.
- [30] B. Ummenhofer, H. Zhou, J. Uhrig, N. Mayer, E. Ilg, A. Dosovitskiy, and T. Brox. Demon: Depth and motion network for learning monocular stereo. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [31] J. Xiao, A. Owens, and A. Torralba. Sun3d: A database of big spaces reconstructed using sfm and object labels. In *Proc. of Int’l Conf. on Computer Vision (ICCV)*, 2013.
- [32] F. Yu and D. Gallup. 3d reconstruction from accidental motion. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [33] J. Zbontar and Y. LeCun. Computing the stereo matching cost with a convolutional neural network. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2015.