

DnD: Dense Depth Estimation in Crowded Dynamic Indoor Scenes

Dongki Jung^{* 1} Jaehoon Choi^{* 1,2} Yonghan Lee¹ Deokhwa Kim¹ Changick Kim³
Dinesh Manocha² Donghwan Lee¹
¹NAVER LABS ²University of Maryland ³KAIST

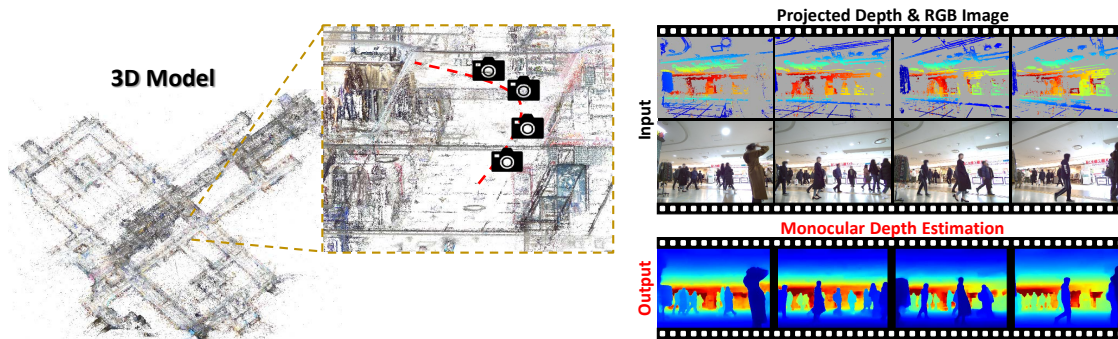


Figure 1: With traditional 3D reconstruction methods [45, 46], we can obtain a 3D model on complex and crowded indoor environments. Our novel approach, DnD, takes both an RGB image and a sparse depth map projected from this 3D model as input. Our method predicts dense and absolute-scale depth maps of single view dynamic scenes.

Abstract

We present a novel approach for estimating depth from a monocular camera as it moves through complex and crowded indoor environments, e.g., a department store or a metro station. Our approach predicts absolute scale depth maps over the entire scene consisting of a static background and multiple moving people, by training on dynamic scenes. Since it is difficult to collect dense depth maps from crowded indoor environments, we design our training framework without requiring depths produced from depth sensing devices. Our network leverages RGB images and sparse depth maps generated from traditional 3D reconstruction methods to estimate dense depth maps. We use two constraints to handle depth for non-rigidly moving people without tracking their motion explicitly. We demonstrate that our approach offers consistent improvements over recent depth estimation methods on the NAVERLABS dataset, which includes complex and crowded scenes.

1. Introduction

There is considerable interest in using robotics and aug-

^{*} These two authors contributed equally.

Correspondence to dongki.jung@naverlabs.com, kevchoi@umd.edu

mented reality technologies in crowded real-world spaces corresponding to malls, airports, or public places. In order to perform safe navigation or combine real and virtual worlds, robots [29, 28, 44] or mobile devices [50, 35] need a 3D geometric representation of large-scale indoor environments. While there is considerable progress in terms of capturing depth using LiDARs or stereo cameras, existing devices still have their own limitations. For example, 3D LiDARs [8] tend to produce sparse depth maps for distant objects and may result in noisy point cloud maps due to the high level of occlusions caused by multiple moving people. Moreover, because of the high prices and large volumes of the depth sensors, there is a critical need to consider the case where only a single camera is available.

Given a large number of video frames, traditional 3D reconstruction methods such as structure-from-motion (SfM) and multi-view stereo (MVS) [14, 49, 45, 46] can generate a 3D model. Recently, visual localization techniques [38] have been developed to allow mobile devices to obtain the location and camera pose from which the image is taken. Given the 3D model and visual position, however, mobile devices moving through crowded indoor environments are only able to capture sparse and highly noisy depth maps. This is because traditional reconstruction methods are based on both a static scene assumption (that static areas of the

scenes can be observed from two different viewpoints) and correct feature matching in two or more images. However, the moving pedestrians in crowded indoor environments tend to violate the static scenes assumption. Moreover, traditional 3D reconstruction methods may fail to perform the correct matching on non-textured (e.g., walls), specular, and reflective regions (e.g., glass) of the scene. Consequently, these problems lead to a 3D model on complex and crowded indoor environments.

Main Results: To address these limitations, we explore new methods that can utilize the 3D model generated using traditional 3D reconstruction methods [45, 46] into learning-based depth estimation algorithms for dynamic scenes. Given the 3D model, our method can be used for general applications because it can compute dense depth maps of dynamic scenes without reconstructing this 3D model iteratively. In contrast to supervised learning methods [27, 5, 26], our method does not rely on dense depth maps generated from depth sensing devices. Our approach, shown in Fig. 1, takes an RGB image and a sparse depth map projected from the 3D model as input and outputs a dense depth map. Given the pose obtained from the SfM [45], we propose using the photometric consistency loss, which enables our method to estimate dense depth maps, and depth loss, forcing our network to learn absolute-scale depth. Although these loss functions are useful for providing dense depth maps in static background regions, there still exist great challenges in estimating depths of multiple non-rigidly moving objects, i.e. the pedestrians. To overcome this limitation, we propose two constraints: 1) a flow-guided shape constraint to refine the depth maps for human regions by filling missing parts of the human regions and removing visual artifacts, and 2) a normal-guided scale constraint to force our neural network to learn the absolute scale depth in human regions guided by depth values in the human’s ground contact point.

Compared to traditional reconstruction approaches or recent learning-based methods [26, 66, 25], our approach (DnD) shows a better ability to predict plausible depth in both human and non-human regions, though we only use a monocular camera. We evaluate our approach on the NAVERLABS dataset [24], the first dataset that provides both metric 3D SfM models and dynamic scenes collected from a department store and a metro station. The main contributions of this paper are summarized as follows:

- We introduce a novel approach to estimate the depth maps using both dynamic scenes collected from a moving monocular camera and given sparse depth maps. We train the monocular depth estimation network with a 3D model generated by traditional 3D reconstruction algorithms [45, 46].
- We present two novel constraints based on optical flow

and surface normal that improve the accuracy of our monocular depth estimation network to predict absolute scale depths for moving people.

- We highlight the benefits of our approach over the state-of-the-art methods on crowded indoor environments and observe 3.6% -10.2% improvement in RMSE. Furthermore, our method works well in diverse indoor datasets like TUM RGB-D and NYUv2.

2. Related Work

Structure from Motion and Multi-view Stereo Traditional SfM systems [49, 1, 45] rely on matching features across two or more images of the same scene and using epipolar geometry [19] to reconstruct depth. Multi-view stereo (MVS) algorithms estimate the dense 3D structure of a scene with multiple calibrated images from the arbitrary viewpoints [14]. Recently, some researchers have presented learning-based MVS methods [62, 66] that build on a neural network to learn the regularization of 3D cost volume. However, traditional SfM and MVS methods often produce sparse and erroneous 3D reconstruction due to incorrect feature matches and dynamic objects [37]. Due to static scene assumption, they either drop pixels with low confidence or estimate incorrect depth values for dynamic objects.

Monocular Depth Estimation The existing learning-based depth estimation methods fall into two groups. One group is based on supervised learning [11, 10, 32, 23, 13, 25, 63], which requires a large-scale dataset with groundtruth depth maps. However, all of these methods require dense groundtruth collected from active depth sensors. Other recent works explore the idea of training with weak supervisions, e.g., using ordinal depth relations as groundtruth [69, 4, 56, 57] and leveraging multi-view stereo reconstruction algorithms to generate pseudo groundtruth from internet photo collections [27, 5]. The other group is based on self-supervised learning using either rectified stereo image pairs [15, 16] or monocular video sequences [68, 64, 17, 18, 59, 61, 6] as training data. Video-based depth estimation methods [31, 53, 34] use temporally consecutive frames to estimate depth over time during inference under the assumption of handling only static scenes. Several works for monocular depth completion [12, 36, 58, 60, 7] have been proposed to capitalize on sparse depth maps with corresponding images, resulting in dense depth estimations.

Depth Estimation of a Dynamic Scene Existing works [43, 41] use object-level motion segmentation to reconstruct a dynamic scene. The key challenge of applying deep neural networks to this task is the lack of large-scale datasets containing diverse and dynamic scenes. Many works adopt a data-driven approach by building diverse datasets from either internet stereo videos [54] or 3D videos [40]. Some works [26, 5] use SfM and MVS to build a dataset with

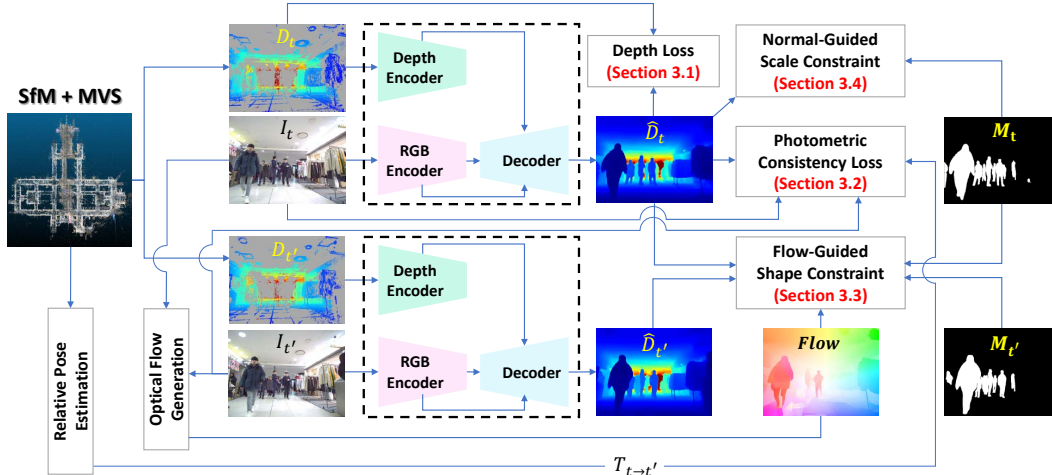


Figure 2: Overview of our proposed approach. The black box with a dotted line shows our monocular depth estimation network consisting of an RGB encoder, a depth encoder, and a decoder. For training, the network takes consecutive temporal frames ($I_t, I_{t'}$) with corresponding sparse depth maps ($D_t, D_{t'}$) projected from the 3D model as input. Our training method is a combination of four terms. 1) The depth loss encourages our network to encode absolute scale from a depth input. 2) The photometric consistency loss is based on view synthesis and regularizes the network training for static background regions. 3) The flow-guided shape constraint enables the network to complete missing pixels in human regions with proper depth values and eliminate visual artifacts. 4) The normal-guided scale constraint enables our network to estimate accurate and absolute-scale depths on moving people.

depth groundtruth from internet video collections. Li et al. [26] use a collection of Mannequin Challenge videos to train a network for depth estimation in dynamic scenes. Yoon et al. [65] present a view synthesis of a dynamic scene via monocular depth estimation with a pre-trained model [40] and MVS. Luo et al. [34] finetune a pre-trained model to satisfy 3D geometric constraints on consecutive video frames. Although their method can handle scenes with a moderate object motion, it is vulnerable to crowded scenes with extreme object motion. In comparison, our method does not require either particular datasets or the pre-trained depth estimation network. Our approach can be generalized to large-scale datasets in real-world crowded indoor environments.

3. Our Method: DnD

We present a learning-based approach to estimate dense and absolute-scale depth maps in scenes from complex and crowded indoor environments. Intuitively, most scenes consist of a static background (e.g., a wall) and dynamic objects (e.g., moving people). The sparse depth maps projected from the 3D model have absolute-scale depth values in small regions of the static background. The rest of the regions in the projected depth maps, including the static background and dynamic objects, have empty depth values. Our proposed approach trains a dense depth estimation network from monocular video sequences with the corresponding sparse depth maps. As shown in Fig. 2, our depth estimation model takes a current image I_t with a sparse depth map

D_t and temporally adjacent images $I_{t'}$ with sparse depth maps $D_{t'}$. The sparse depth maps D_t and $D_{t'}$ are projected from the 3D model. The $I_{t'}$ includes two temporal frames I_{t-1} and I_{t+1} . Our depth estimation model predicts dense depth maps \hat{D}_t and $\hat{D}_{t'}$ to compute a flow-guided shape constraint.

3.1. Absolute-Scale Depth Loss

Generally, monocular depth estimation methods suffer from an inherent scale ambiguity problem. To mitigate this limitation, we use an absolute-scale depth input D_t as a groundtruth. We apply the L1 loss to penalize the differences between the depth input D_t and the depth prediction \hat{D}_t on the pixels where D_t values exist. The depth loss is formulated as,

$$L_d = \sum_{p \in \Omega} \| \hat{D}_t(p) - D_t(p) \|, \quad (1)$$

where Ω indicates valid points that involve available sparse depths. This loss enables our network to learn absolute-scale depth values in the regions of the depth input and further extrapolate depth values with absolute scales in empty regions including a static background and moving people.

3.2. Photometric Consistency Loss

To train our network, we turn to previous self-supervised monocular depth estimation methods [68, 17] that have utilized the photometric loss as the main loss function for training the network. We are able to generate a synthesized

frame I'_t by reprojecting temporally adjacent frames $I_{t'}$ to the current frame I_t with the camera intrinsic matrix K , the predicted depth \hat{D}_t , and the relative pose $T_{t \rightarrow t'}$. This relative pose between consecutive temporal frames is computed from the absolute poses from SfM [45]. The photometric consistency loss with a combination of L1 and SSIM [55] is formulated as follows:

$$L_{ph} = \alpha \frac{1 - \text{SSIM}(I_t, I'_t)}{2} + (1 - \alpha) \| I_t - I'_t \|, \quad (2)$$

$$I'_t(p) = I_{t'} \langle \pi(KT_{t \rightarrow t'} \hat{D}_t(p) K^{-1} \tilde{p}) \rangle,$$

where $\alpha = 0.85$, \tilde{p} is the homogeneous coordinate of p , π means projection from homogeneous to image coordinates, and $\langle \cdot \rangle$ indicates the bilinear sampling function. Note that one current frame and two temporally adjacent frames are used to compute the photometric consistency loss. Following [17], we remove the occluded and out-of-view pixels, and also mask out low texture regions via minimum re-projection loss and auto-masking techniques. In particular, this photometric consistency loss encourages our network to predict dense depth maps over the static backgrounds.

3.3. Flow-Guided Shape Constraint

The loss functions described in the previous sections are still limited in their ability to provide accurate depths when there are moving people. This is because both MVS and photometric consistency loss operate under the assumption of a single moving camera and a static scene. This assumption implies that inconsistencies between different views of each image are only derived from the camera ego-motion in static scenes [19]. If we correctly estimate the motion of moving people between the consecutive frames, the view consistency for 3D reconstruction can be obtained for two different views. However, with multiple moving people in dynamic scenes, triangulation-based methods fail to achieve a scene consistent depth map. Since multiple moving people in scenes undergo non-rigid deformations, it is difficult to explicitly estimate their 3D motion [54, 26, 22].

Instead of modeling object motions in 3D, we leverage monocular video frames to estimate temporally coherent depth maps with respect to moving camera motions and non-rigidly moving people. To achieve this, we use optical flow F , a human mask M_t from the current frame I_t , and human masks $M_{t'}$ from the temporally adjacent frames $I_{t'}$. Given consecutive frame pairs, the optical flow describes which pixel pairs have the same intensities. Human masks are also used to find correct regions that overlap due to the high-level of occlusion, non-rigid deformation, and complex ego-motions of the moving camera.

$$M = M_t \cap F_{t' \rightarrow t}(M_{t'}), \quad (3)$$

where M is the overlapped region of the flow-warped human mask $M_{t'}$ and the current human mask M_t to prevent

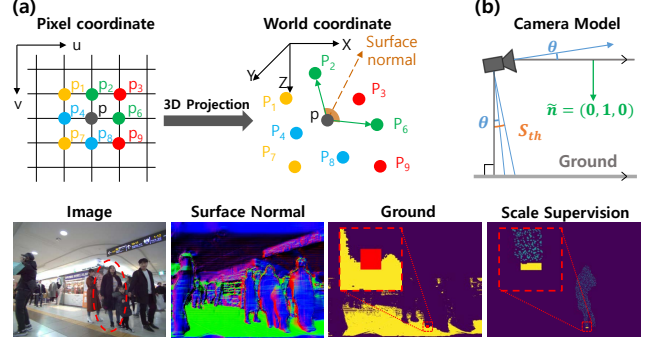


Figure 3: (a) The process to obtain a surface normal from the pixel coordinates and the corresponding depth values. (b) shows that the ground normal \tilde{n} is not aligned with the y-direction of the camera coordinates. The bottom images indicate the process for propagating the reliable scale for the person in the dashed red circle. The red point in the third column means the small patch B , indicating the human’s ground contact point. In the fourth column, the yellow and blue points show the ground-patch intersection and uniformly sampled points in the human mask respectively.

occlusion issues, meaning that some pixels from M_t do not appear in $F_{t' \rightarrow t}(M_{t'})$. In dynamic scenes captured by the moving camera, the absolute depth values for dynamic human regions can be inconsistent. Thus, we apply the optical flow field to compare consecutive inverse depth predictions for human regions in gradient domains. Our proposed flow-guided shape constraint is defined as

$$\bar{\nabla}(d^*(p)) = \frac{\nabla d^*(p)}{|d^*(p) + \nabla d^*(p)| + |d^*(p)|}, \quad (4)$$

$$L_f = \frac{1}{|M|} \sum_{p \in M} |\bar{\nabla} d_t^*(p) - \bar{\nabla} F_{t' \rightarrow t}(d_{t'}^*(p))|$$

where ∇ and $\bar{\nabla}$ denote the gradient and scale-invariant gradient and $d_t^* = d_t / \mu(d_t)$ indicates the mean-normalized inverse depth. Our flow-guided shape constraint enforces smooth gradients and completes missing parts of human regions with accurate depth values. To regularize the depth in a static background, we impose the edge-aware depth smoothness loss L_s over pixels in non-human regions M^C :

$$L_s = \sum_{p \in M^C} |\nabla d_t^*(p)| e^{-|\nabla I_t(p)|}. \quad (5)$$

3.4. Normal-Guided Scale Constraint

To achieve scale-aware depth estimation, accurate depth estimation with an absolute scale for moving people is crucial. We leverage absolute-scale depth values in the static background to constrain the scale of depth values for moving people. Intuitively, people are standing on the ground and the depth values of the humans are nearly the same as

the ones on the ground. The human’s ground contact point can be a significant geometric cue for scale-aware depth estimation. Thus, we introduce weak supervision, which forces several randomly selected points corresponding to the particular human mask to be consistent with the ground contact points’ depth values.

Inspired by [61, 59], we estimate the ground area using the surface normal. The 8-neighbors convention is used to specify the normal direction, as seen in Fig. 3 (a). Those 8 points at the 2D coordinate are split into 4 pairs, where each pair of vectors is perpendicular, e.g., $(i, j) = \{(2, 6), (3, 9), \dots\}$. The average of the four normals determines the final surface normal for the position p :

$$N_t(p) = \frac{1}{4} \sum_{\vec{pp}_i \perp \vec{pp}_j} \frac{\vec{PP}_i \times \vec{PP}_j}{\|\vec{PP}_i \times \vec{PP}_j\|_2}, \quad (6)$$

where P means the reconstructed 3D point clouds for pixel p by the predicted depth and the operator \times denotes the cross product. Among many pixel points p , we only need points in ground regions. The true normal direction of ground $\tilde{n} = (0, 1, 0)$ and the estimated normal $N_t(p)$ should be matched. As in Fig. 3 (b), the camera is not perpendicular to the ground by a difference of θ for the camera model, and there exists the uncertainty of the estimated normal $N_t(p)$. Thus, we set S_{th} as a threshold to find the ground area and calculate the cosine similarity between the true ground normal \tilde{n} and the predicted normal $N_t(p)$:

$$G = \{p \mid |\cos^{-1} \frac{\tilde{n} \cdot N_t(p)}{\|\tilde{n}\| \|N_t(p)\|}| < S_{th}\}, \quad (7)$$

where the operator \cdot denotes the inner product. Then, in order to determine the human’s ground contact point, we detect small patches B centered at the bottom of the pixel in the human instances. These pixels in both the small patches and the ground area are possible candidates for the human’s ground contact point. However, simply applying the constant value as supervision risks severe degradation of depth values in the human region because moving people have exquisite and isometric shapes. Therefore, we uniformly sample the pixels M' in the human mask M_t and constrain them by the median value in the intersection of B and G :

$$L_n = \sum_{p \in M'} \frac{|\hat{D}_t(p) - med(\hat{D}_t(B \cap G))|}{\hat{D}_t(p)}, \quad (8)$$

where the operator med takes the median value of factors.

3.5. Overall Loss Function

Our overall loss function is a weighted sum over the previously introduced losses,

$$L_{total} = \lambda_d L_d + \lambda_p L_p + \lambda_s L_s + \lambda_f L_f + \lambda_n L_n \quad (9)$$

where $\lambda_d, \lambda_p, \lambda_s, \lambda_f$ and λ_n denote weights on the respective loss terms selected through a grid search.

4. Experiments

4.1. Experiment Settings

To validate our method, we experiment with the NAVERLABS Dataset, NYUv2 and TUM RGB-D dataset. The depth value is represented in the metric scale (m). We evaluate our method using the standard metrics [11].

NAVERLABS Indoor Localization: The NAVERLABS dataset ¹ [24] comprises scenes collected from two different places: Department Store (Dept) and Metro Station (MS). This dataset is split into 60K images for training and 835 images for testing in MS, and 25K images for training and 443 images for testing in Dept. To evaluate the crowdedness of indoor datasets, we use a crowd density, which is a ratio between the area of human pixels in scenes and the area of all image pixels. Most public indoor datasets [48, 2, 9, 3, 47, 52] do not have scenes with moving people ($< 0.1\%$ crowd density). In NAVERLABS, Dept and MS have 6.87% and 12.9% crowd density, the most crowded dataset. Additionally, Dept and MS have 6.7 and 3.6 people per scene on average. Different types of datasets [4, 56, 54, 26] that include diverse dynamic scenes contain only images or ordinal depth relations for pairs of points.

NAVERLABS used all the images collected from 6 cameras to build a 3D model via COLMAP [45, 46]. The groundtruth poses for the input images were initially calculated by LiDAR SLAM and further refined by bundle adjustment with prior results. Afterwards, metric depth images were generated by the COLMAP MVS algorithm. For evaluation, we accumulated 0.3 seconds of LiDAR scans corresponding to one scan line since the LiDAR is sampled at 10 Hz. It is difficult to accumulate more LiDAR sweeps because pedestrians cause severe noise and artifacts in the projected depth maps.

TUM RGB-D: The TUM RGB-D dataset [51] is comprised of 39 video sequences recorded with a Microsoft Kinect. We used a subset of sequence, *walking*, which contains moderate dynamic scenes that two people walk both in the background and the foreground. We train our method on these subsets and evaluate on the test set following MC [26].

NYUv2: Although the NYUv2 dataset [48] does not contain dynamic scenes, it is the most common public benchmark of depth estimation for indoor scenes. The training dataset consists of 268K images. We subsampled every 10 frames and used only 47k frames for training. The official 654 test scenes are used for evaluation.

Implementation Details: Our monocular depth estimation network is based on U-Net [42], an encoder-decoder architecture. For training, we apply FlowNet2.0 [21] to compute optical flow and Mask R-CNN [20] pre-trained on the COCO dataset [30] to detect human masks. For normal-guided scale constraint, we apply 20×20 boxes to compose

¹is available at <https://www.naverlabs.com/datasets>

Method	Data	Scaling	F+B / F (Lower is better)				F+B / F (Higher is better)		
			Abs Rel	Sq Rel	RMSE	RMSE log	$\delta_{1.25}$	$\delta_{1.25^2}$	$\delta_{1.25^3}$
Fast-MVSNet [66]	MS	-	0.383 / 0.754	0.60 / 1.37	2.05 / 2.95	0.123 / 0.213	0.634 / 0.358	0.787 / 0.547	0.866 / 0.697
DepthComple [36]	MS	-	0.800 / 0.863	0.7 / 0.79	4.52 / 4.67	- / -	0.041 / 0.024	0.087 / 0.050	0.139 / 0.083
MiDaS [40]	MS	median	0.413 / 0.459	0.33 / 0.42	2.77 / 2.86	0.181 / 0.191	0.380 / 0.352	0.627 / 0.604	0.785 / 0.773
MC [26]	MS	median	0.294 / 0.315	0.30 / 0.22	2.05 / 2.53	0.122 / 0.162	0.585 / 0.459	0.822 / 0.741	0.906 / 0.847
BTS [25]	MS	-	0.396 / 0.696	0.66 / 1.24	1.96 / 2.70	0.120 / 0.198	0.665 / 0.401	0.792 / 0.589	0.862 / 0.726
DnD (L_{ph} only)	MS	median	0.326 / 0.596	1.38 / 4.16	2.39 / 3.60	0.107 / 0.177	0.711 / 0.545	0.840 / 0.705	0.898 / 0.795
DnD	MS	-	0.189 / 0.240	0.20 / 0.16	1.76 / 2.44	0.084 / 0.133	0.806 / 0.677	0.881 / 0.798	0.919 / 0.856
Fast-MVSNet [66]	Dept	-	0.551 / 0.889	0.98 / 1.97	3.24 / 4.11	- / 0.243	0.431 / 0.288	0.655 / 0.498	0.791 / 0.654
DepthComple [36]	Dept	-	0.842 / 0.897	0.76 / 0.84	5.88 / 6.65	- / -	0.014 / 0.058	0.028 / 0.028	0.093 / 0.046
MiDaS [40]	Dept	median	0.461 / 0.519	0.46 / 0.54	4.39 / 3.72	0.208 / 0.210	0.375 / 0.332	0.579 / 0.554	0.729 / 0.728
MC [26]	Dept	median	0.428 / 0.385	0.40 / 0.29	3.91 / 3.36	0.190 / 0.178	0.364 / 0.384	0.626 / 0.676	0.786 / 0.822
BTS [25]	Dept	-	0.584 / 1.066	1.36 / 2.80	3.06 / 4.40	0.159 / 0.260	0.561 / 0.327	0.721 / 0.488	0.808 / 0.614
DnD (L_{ph} only)	Dept	median	0.289 / 0.388	0.47 / 0.60	2.60 / 3.24	0.109 / 0.148	0.663 / 0.564	0.837 / 0.759	0.911 / 0.847
DnD	Dept	-	0.213 / 0.250	0.32 / 0.30	2.36 / 3.04	0.084 / 0.116	0.761 / 0.707	0.889 / 0.836	0.932 / 0.886

Table 1: Quantitative comparisons with the state-of-the-art depth estimation algorithms. F means the evaluation results in the human regions and F+B indicates the evaluation results on depth values over the entire scene. MS and Dept denote the Metro Station dataset and the Department Store dataset, respectively. DnD (L_{ph} only) is trained only using the photometric consistency loss. In the MS dataset, DnD shows improvement in depth for human regions by 3.6% in the RMSE metric. In terms of depth for entire regions, DnD reduces the RMSE by 10.2%.

the small patch B . The camera of the NAVERLABS dataset is towards the upper side for $\theta = 10^\circ$ as in Fig. 3 (b); thus S_{th} is set to 15° . We investigate the sampling ratio for M' in Eq. 8 at 10%, 30%, and 50%, and we select 30% because it is slightly better than others. In the supplementary material, we provide more implementation details of network architecture and training procedures.

4.2. Experimental Results on NAVERLABS dataset

For quantitative evaluation, we compare our proposed method DnD with five different methods: 1) Fast-MVSNet [66] is a learning-based multiview stereo algorithm, 2) DepthComple is a learning-based depth completion algorithm that adopts an early-fusion encoder-decoder network from [36] combined with normalized convolution layers [12], 3) MiDaS [40] is a monocular depth estimation method trained on a diverse set of datasets including various dynamic scenes, 4) MC [26] is a monocular depth estimation method trained on a Mannequin Challenge dataset, and 5) BTS [25] is the state-of-the-art supervised monocular depth estimation method on the benchmark datasets [48]. To experiment with Fast-MVSNet, DepthComple, and BTS, we use their public codes and train the network from scratch with the NAVERLABS dataset. MiDaS and MC are implemented with pre-trained weights from their public codes. To train the former methods, we exploit a depth map projected from the 3D model as groundtruth. In our experiments, training video-based depth estimation methods [34, 53] fails in the video frames of NAVERLABS. Their methods based on static scene assumption cannot be applied for scenes with extreme object motion. Furthermore, self-supervised methods [17, 36] show poor performance for depth evaluation. We observe that a joint training framework of pose and depth from a monocular video is extremely difficult be-

cause pose networks often fail to estimate proper camera ego-motion in complex and crowded indoor environments. Instead, we provide experimental results of DnD (L_{ph} only) trained only with the photometric consistency loss similar to self-supervised training settings, e.g., Monodepth2 [17]. DnD (L_{ph} only) only converts an RGB image to a dense depth map and thus suffers from the scale ambiguity issue. DnD, Fast-MVSNet, DepthComple, and BTS are able to estimate absolute depth values while MiDaS and MC only output relative depth up to an unknown scale factor. Thus, we apply a median scaling [17] to define a scale factor by comparing their depth predictions with the absolute depth from our groundtruth depths. In Table 1, we summarize the quantitative results on two indoor datasets. DnD outperforms other depth estimation methods on all evaluation metrics. In particular, we observe an improvement in performance at depths of both people (F) and the entire scene (F+B). Although the models provided by MiDaS and MC were trained on different datasets, the methods focus on composing datasets for dynamic scenes and training their depth estimation network with supervised learning. Compared to such methods, our method can be scalable to general environments without high quality groundtruth depths.

A qualitative comparison is shown in Fig. 4. The visual results of DepthComple are added in the supplementary material because this method fails to generate reasonable depth maps. In general, the qualitative comparison corresponds well to our quantitative results in Table 1. In terms of depth in human regions, Fast-MVSNet and BTS fail to reconstruct the depth for moving people regions. BTS is not able to learn depths for moving people regions because it only relies on COLMAP results as ground truth. Fast-MVSNet is also not able to provide depths for moving people due to the limitation of multi-view matching. In contrast, MC and Mi-

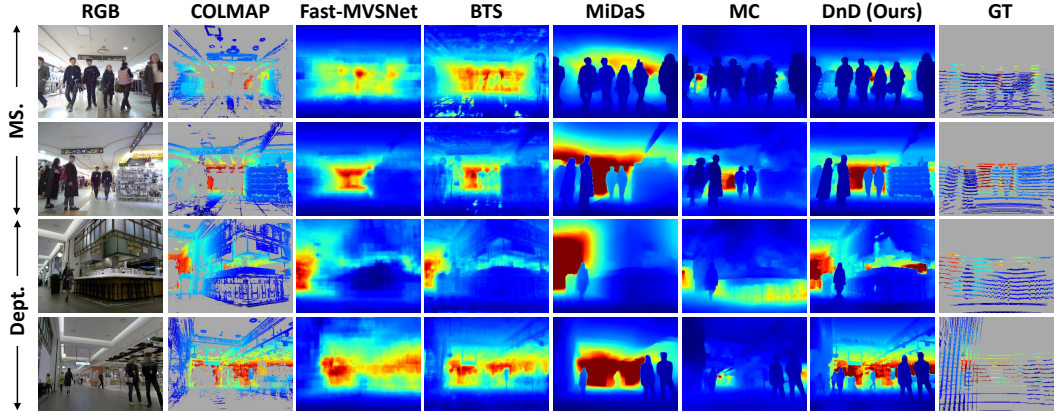



Figure 4: Qualitative comparisons of the Metro Station dataset (1-2 rows) and the Department Store dataset (3-4 rows). Our network takes an image (1st column) and a depth input projected from the 3D model (2nd column), and predicts dense depth maps (7th column). Ground truth (8th column) collected from 3D LiDAR is used to evaluate our method. Depth results (3-6 columns) are estimated by the existing depth estimation algorithms. Compared to MiDaS and MC, DnD generates detailed and high quality depth maps in non-human regions and also shows depth maps in human regions with sharp depth discontinuities. All color maps use the jet color map (low  high; grey means empty depth values).

Method	F+B / F (Lower is better)				F+B / F (Higher is better)		
	Abs Rel	Sq Rel	RMSE	RMSE log	$\delta_{1.25}$	$\delta_{1.25^2}$	$\delta_{1.25^3}$
DnD (L_{ph} only)	0.326 / 0.596	1.38 / 4.16	2.39 / 3.60	0.107 / 0.177	0.711 / 0.545	0.840 / 0.705	0.898 / 0.795
DnD w/o FSC, NSC	0.209 / 0.304	0.24 / 0.25	1.80 / 2.56	0.095 / 0.172	0.767 / 0.540	0.851 / 0.690	0.899 / 0.787
DnD w/o NSC	0.204 / 0.299	0.23 / 0.24	1.78 / 2.51	0.094 / 0.170	0.772 / 0.552	0.856 / 0.701	0.901 / 0.790
DnD w/o FSC	0.192 / 0.252	0.20 / 0.17	1.77 / 2.47	0.086 / 0.139	0.794 / 0.644	0.877 / 0.785	0.917 / 0.851
DnD (full)	0.189 / 0.240	0.20 / 0.16	1.76 / 2.44	0.084 / 0.133	0.806 / 0.677	0.881 / 0.798	0.919 / 0.856

Table 2: Contributions of our proposed modules to the evaluation results on the Metro Station (MS) dataset. F means the evaluation results on depth values in the human regions and F+B indicates the evaluation results on depth values over the entire scene. The median scaling is applied to DnD (L_{ph} only) for absolute scale depth prediction. Compared to DnD (L_{ph} only), we improve the RMSE by 26.4% in the entire scene and by 32.2% in the human regions.

DaS adequately estimate the depths in human regions with sharp depth discontinuities. However, these methods show blurry and noisy depths for background regions. Compared to other methods, depth predictions over background regions from our method are sharper and less noisy.

Ablation Study In Table 2, we validate the influence of different loss terms within our proposed method. We note that all our proposed modules contribute to a significant improvement in Table 2. In particular, both NSC and FSC significantly improve depth estimation performance for all evaluation metrics. Figure 5 provides an ablation study using qualitative comparison to better understand the effectiveness of different loss terms in our method. In Fig. 5 (a), only using depth loss and photometric consistency loss for training is able to estimate dense depth maps over the entire scene. However, they show poor results especially on depth in human regions. Figure 5 (b) and (c) validate that our method improves the depth estimation performance with each proposed module.

Analysis We study the effectiveness of the number of sparse depth points and the crowd density. We use the accuracy

score $\delta_{1.25}$ as an evaluation metric because other metrics do not show significant performance degradation in DnD. If the number of matched local features is not enough for reconstructing a plausible 3D model, the projected depth input for DnD might become a highly sparse map. In this case, the experimental results in Fig. 6 (a) show the relationship between the number of depth points and the performance of DnD. As the number of input depth points decreases, the performance degrades slightly as expected. However, this robustness to the sparsity proves that the RGB images serve as the primary source of our monocular depth estimation network. Furthermore, for the dense crowded scenes in indoor environments, we select some highly crowded samples in the original test set for the NAVERLABS dataset. These are composed of 12% or more human pixels out of the total number of the pixels in one image. Note that the crowd density is a more significant factor than the number of people per scene in our task. We focus on pedestrians whose apparent sizes are large (i.e., high crowd density) because they are close to the camera and occlude those who lie behind them. Despite the multiple non-rigid people, we notice that DnD

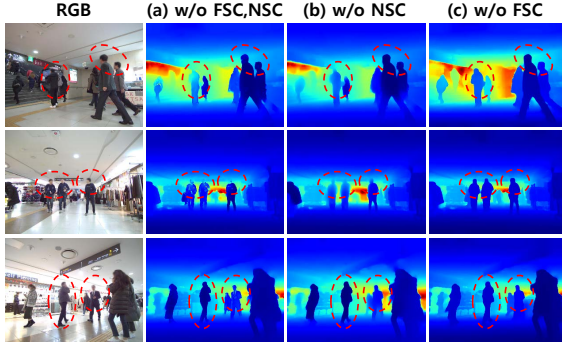


Figure 5: Qualitative comparisons of the proposed modules. (a) We note that a part of the human is cut off in the depth map. The visual artifacts that look similar to RGB image patterns often appear in human depth regions. (b) FSC encourages our network to capture sharper details, fill in missing parts of the human regions, and remove visual artifacts. (c) We observe that NSC estimates depth with sharp depth discontinuities and also recovers absolute scale depth values for human regions.

robustly estimates depth for both humans and background regions as in Fig. 6 (b).

DnD with Localization For the complete system of DnD, there is a need for reprojected depth maps obtained from the 3D model using visual localization techniques. Thus, we exploit Kapture toolbox [39] to estimate the camera pose of given input images, and then project the 3D model to the sparse depth maps. We describe more details about the whole pipeline in the supplementary materials.

4.3. Results on NYUv2 and TUM RGB-D datasets

We train and evaluate our method on each NYUv2 and TUM RGB-D dataset. Since the 3D map for TUM and NYUv2 produced by COLMAP is only up to an unknown scale factor, we cannot use metric depth as input. Instead, we obtain the metric depth maps by sparsifying depth maps collected from Kinect in order to use standard error metrics for a fair comparison. To simulate sparse patterns existing in COLMAP, we only maintain the depth values according to the location of SIFT features [33] in the corresponding images. Since the number of SIFT features varies in each frame, we set the maximum number of features to 200 samples. Table 3 reports the quantitative comparisons with several other depth estimation methods on both datasets. DnD outperforms other single view methods, demonstrating the benefit of our method in less crowded benchmark indoor datasets. In supplementary materials, we show experimental results on TUM when we consider normalized relative depth as projected depth maps obtained from a scale ambiguous 3D model.




Figure 6: (a) shows $\delta_{1.25}$ metric for our method corresponding the number of uniform-sampled input depth points, and (b) indicates the performance degradation for scenes with high density crowds in MS and Dept: 8.8% and 9.7% for FMVS (Fast-MVSNet [66]), 20% and 5.2% for MC [26], 10.2% and 17.1% for BTS [25], 6.2% and 6% for ours.

Method	Dataset	Lower is better		Higher is better		
		RMSE	Abs Rel	$\delta_{1.25}$	$\delta_{1.25^2}$	$\delta_{1.25^3}$
BTS [25] (Single view)	NYUv2	0.392	0.110	0.885	0.978	0.994
Yang et al. [60] (Single view)	NYUv2	0.569	0.171	-	-	-
MonoDepth2 [17] (Single view)	NYUv2	0.617	0.170	0.748	0.942	0.986
P2Net [67] (Multi view)	NYUv2	0.533	0.147	0.801	0.951	0.987
DeepV2D [53] (Multi view)	NYUv2	0.403	0.061	0.956	0.988	0.996
DnD (Single view)	NYUv2	0.362	0.098	0.910	0.988	0.998
MC [26] (Single view)	TUM	0.840	0.204	0.664	0.931	0.981
MC [26] (Multi view)	TUM	0.570	0.129	-	-	-
MiDaS [40] (Single view)	TUM	0.819	0.196	0.678	0.911	0.965
DnD (Single view)	TUM	0.631	0.175	0.751	0.947	0.982

Table 3: Quantitative comparisons of the NYUv2 and TUM RGB-D dataset. In both static and dynamic scenes, DnD shows improved performance compared to recent methods.

5. Conclusion, Limitations, and Future Work

In this paper, we present a learning-based method for estimating dense depth maps of dynamic scenes collected from a moving monocular camera. Given RGB images and sparse depth maps projected from a 3D model, our method is able to predict absolute scale depth for multiple pedestrians and complex backgrounds. In complex and crowded indoor environments, this is a practical method to use for various robotics and augmented reality applications. Once we initially have the 3D model, our method enables any freely moving mobile devices with a single-view camera to estimate dense depth maps. Our method still has limitations that we plan to address. The sparse depth maps projected from the 3D model may be inaccurate or not aligned with the current image. Also, our training method relies on FlowNet2.0 [21] and Mask R-CNN [20]. In the future, we would like to extend our approach to outdoor scenes.

Acknowledgements This work was supported by the Institute of Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. 2019-0-01309, Development of AI Technology for Guidance of a Mobile Robot to its Goal with Uncertain Maps in Indoor/Outdoor Environments)

References

- [1] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54(10):105–112, 2011. [2](#)
- [2] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017. [5](#)
- [3] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017. [5](#)
- [4] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. In *Advances in neural information processing systems*, pages 730–738, 2016. [2](#), [5](#)
- [5] Weifeng Chen, Shengyi Qian, and Jia Deng. Learning single-image depth from videos using quality assessment networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5604–5613, 2019. [2](#)
- [6] Jaehoon Choi, Dongki Jung, Donghwan Lee, and Changick Kim. Safenet: Self-supervised monocular depth estimation with semantic-aware feature extraction. *arXiv preprint arXiv:2010.02893*, 2020. [2](#)
- [7] Jaehoon Choi, Dongki Jung, Yonghan Lee, Deokhwa Kim, Dinesh Manocha, and Donghwan Lee. Selfdeco: Self-supervised monocular depth completion in challenging indoor environments. In *Proc. of the IEEE International Conference on Robotics and Automation*, Xi’an, China, 2021. [2](#)
- [8] John A Christian and Scott Cryan. A survey of lidar technology and its use in spacecraft relative navigation. In *AIAA Guidance, Navigation, and Control (GNC) Conference*, page 4641, 2013. [1](#)
- [9] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5828–5839, 2017. [5](#)
- [10] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE international conference on computer vision*, pages 2650–2658, 2015. [2](#)
- [11] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, pages 2366–2374, 2014. [2](#), [5](#)
- [12] Abdelrahman Eldesokey, Michael Felsberg, and Fahad Shahbaz Khan. Confidence propagation through cnns for guided sparse depth regression. *IEEE transactions on pattern analysis and machine intelligence*, 2019. [2](#), [6](#)
- [13] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2002–2011, 2018. [2](#)
- [14] Yasutaka Furukawa and Carlos Hernández. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015. [1](#), [2](#)
- [15] Ravi Garg, Vijay Kumar BG, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *European Conference on Computer Vision*, pages 740–756. Springer, 2016. [2](#)
- [16] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 270–279, 2017. [2](#)
- [17] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 3828–3838, 2019. [2](#), [3](#), [4](#), [6](#), [8](#)
- [18] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2485–2494, 2020. [2](#)
- [19] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. [2](#), [4](#)
- [20] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. [5](#), [8](#)
- [21] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017. [5](#), [8](#)
- [22] Suryansh Kumar, Ram Srivatsav Ghorakavi, Yuchao Dai, and Hongdong Li. Dense depth estimation of a complex dynamic scene without explicit 3d motion estimation. *arXiv preprint arXiv:1902.03791*, 2019. [4](#)
- [23] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *2016 Fourth international conference on 3D vision (3DV)*, pages 239–248. IEEE, 2016. [2](#)
- [24] Donghwan Lee, Soohyun Ryu, Suyong Yeon, Yonghan Lee, Deokhwa Kim, Cheolho Han, Yohann Cabon, Philippe Weinzaepfel, Nicolas Guérin, Gabriela Csurka, et al. Large-scale localization datasets in crowded indoor spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3227–3236, 2021. [2](#), [5](#)
- [25] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019. [2](#), [6](#), [8](#)
- [26] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T Freeman. Learning the depths of moving people by watching frozen people. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4521–4530, 2019. [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)

- [27] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2041–2050, 2018. 2
- [28] Jing Liang, Utsav Patel, Adarsh Jagan Sathyamoorthy, and Dinesh Manocha. Crowdsteer: Realtime smooth and collision-free robot navigation in dense crowd scenarios trained using high-fidelity simulation. *arXiv preprint arXiv:2004.03089*, 2020. 1
- [29] Jing Liang, Yi-Ling Qiao, and Dinesh Manocha. Of-vo: Reliable navigation among pedestrians using commodity sensors. *arXiv preprint arXiv:2004.10976*, 2020. 1
- [30] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 5
- [31] Chao Liu, Jinwei Gu, Kihwan Kim, Srinivasa G Narasimhan, and Jan Kautz. Neural rgb (r) d sensing: Depth and uncertainty from a video camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10986–10995, 2019. 2
- [32] Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE transactions on pattern analysis and machine intelligence*, 38(10):2024–2039, 2015. 2
- [33] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 2004. 8
- [34] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Transactions on Graphics (TOG)*, 2020. 2, 3, 6
- [35] Simon Lynen, Bernhard Zeisl, Dror Aiger, Michael Bosse, Joel Hesck, Marc Pollefeys, Roland Siegwart, and Torsten Sattler. Large-scale, real-time visual–inertial localization revisited. *The International Journal of Robotics Research*, 39(9):1061–1084, 2020. 1
- [36] Fangchang Ma, Guilherme Venturelli Cavalheiro, and Sertac Karaman. Self-supervised sparse-to-dense: Self-supervised depth completion from lidar and monocular camera. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3288–3295. IEEE, 2019. 2, 6
- [37] Onur Ozyesil, Vladislav Voroninski, Ronen Basri, and Amit Singer. A survey of structure from motion. *arXiv preprint arXiv:1701.08493*, 2017. 2
- [38] Nathan Piasco, Désiré Sidibé, Cédric Demonceaux, and Valérie Gouet-Brunet. A survey on visual-based localization: On the benefit of heterogeneous data. *Pattern Recognition*, 74:90–109, 2018. 1
- [39] Noé Pion, Martin Humenberger, Gabriela Csurka, Johann Cabon, and Torsten Sattler. Benchmarking image retrieval for visual localization. In *2020 International Conference on 3D Vision (3DV)*, pages 483–494. IEEE, 2020. 8
- [40] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 2, 3, 6, 8
- [41] Rene Ranftl, Vibhav Vineet, Qifeng Chen, and Vladlen Koltun. Dense monocular depth estimation in complex dynamic scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4058–4066, 2016. 2
- [42] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 5
- [43] Chris Russell, Rui Yu, and Lourdes Agapito. Video pop-up: Monocular 3d reconstruction of dynamic scenes. In *European conference on computer vision*, pages 583–598. Springer, 2014. 2
- [44] Adarsh Jagan Sathyamoorthy, Jing Liang, Utsav Patel, Tianrui Guan, Rohan Chandra, and Dinesh Manocha. Densecavoid: Real-time navigation in dense crowds using anticipatory behaviors. *arXiv preprint arXiv:2002.03038*, 2020. 1
- [45] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016. 1, 2, 4, 5
- [46] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision*, pages 501–518. Springer, 2016. 1, 2, 5
- [47] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene coordinate regression forests for camera relocation in rgb-d images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2930–2937, 2013. 5
- [48] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European conference on computer vision*, pages 746–760. Springer, 2012. 5, 6
- [49] Noah Snavely, Steven M Seitz, and Richard Szeliski. Modeling the world from internet photo collections. *International journal of computer vision*, 80(2):189–210, 2008. 1, 2
- [50] Gregory J Stein, Christopher Bradley, and Nicholas Roy. Learning over subgoals for efficient navigation of structured, unknown environments. In *Conference on Robot Learning*, pages 213–222, 2018. 1
- [51] Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A benchmark for the evaluation of rgb-d slam systems. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 573–580. IEEE, 2012. 5
- [52] Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, and Akihiko Torii. Inloc: Indoor visual localization with dense matching and view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7199–7209, 2018. 5

- [53] Zachary Teed and Jia Deng. Deepv2d: Video to depth with differentiable structure from motion. In *International Conference on Learning Representations*, 2020. 2, 6, 8
- [54] Chaoyang Wang, Simon Lucey, Federico Perazzi, and Oliver Wang. Web stereo video supervision for depth prediction from dynamic scenes. In *2019 International Conference on 3D Vision (3DV)*, pages 348–357. IEEE, 2019. 2, 4, 5
- [55] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 4
- [56] Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, Yang Xiao, Ruibo Li, and Zhenbo Luo. Monocular relative depth perception with web stereo data supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 311–320, 2018. 2, 5
- [57] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. Structure-guided ranking loss for single image depth prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 611–620, 2020. 2
- [58] Yan Xu, Xinge Zhu, Jianping Shi, Guofeng Zhang, Hujun Bao, and Hongsheng Li. Depth completion from sparse lidar data with depth-normal constraints. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2811–2820, 2019. 2
- [59] Feng Xue, Guirong Zhuo, Ziyuan Huang, Wufei Fu, Zhuoyue Wu, and Marcelo H Ang Jr. Toward hierarchical self-supervised monocular absolute depth estimation for autonomous driving applications. *arXiv preprint arXiv:2004.05560*, 2020. 2, 5
- [60] Yanchao Yang, Alex Wong, and Stefano Soatto. Dense depth posterior (ddp) from single image and sparse range. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3353–3362, 2019. 2, 8
- [61] Z Yang, P Wang, W Xu, L Zhao, and R Nevatia. Unsupervised learning of geometry with edge-aware depth-normal consistency. arxiv 2017. *arXiv preprint arXiv:1711.03665*. 2, 5
- [62] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 767–783, 2018. 2
- [63] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5684–5693, 2019. 2
- [64] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1983–1992, 2018. 2
- [65] Jae Shin Yoon, Kihwan Kim, Orazio Gallo, Hyun Soo Park, and Jan Kautz. Novel view synthesis of dynamic scenes with globally coherent depths from a monocular camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5336–5345, 2020. 3
- [66] Zehao Yu and Shenghua Gao. Fast-mvsnet: Sparse-to-dense multi-view stereo with learned propagation and gauss-newton refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1949–1958, 2020. 2, 6, 8
- [67] Zehao Yu, Lei Jin, and Shenghua Gao. P2net: Patch-match and plane-regularization for unsupervised indoor depth estimation. *arXiv preprint arXiv:2007.07696*, 2020. 8
- [68] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1851–1858, 2017. 2, 3
- [69] Daniel Zoran, Phillip Isola, Dilip Krishnan, and William T Freeman. Learning ordinal relationships for mid-level vision. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 388–396, 2015. 2