

Moving Indoor: Unsupervised Video Depth Learning in Challenging Environments

Junsheng Zhou^{1,*}, Yuwang Wang², Kaihuai Qin¹, and Wenjun Zeng²

¹Tsinghua University, Beijing, China

zhoujs17@mails.tsinghua.edu.cn, qkh-dcs@mail.tsinghua.edu.cn

²Microsoft Research, Beijing, China

yuwwan, wezeng@microsoft.com

Abstract

Recently unsupervised learning of depth from videos has made remarkable progress and the results are comparable to fully supervised methods in outdoor scenes like KITTI. However, there still exist great challenges when directly applying this technology in indoor environments, e.g., large areas of non-texture regions like white wall, more complex ego-motion of handheld camera, transparent glasses and shiny objects. To overcome these problems, we propose a new optical-flow based training paradigm which reduces the difficulty of unsupervised learning by providing a clearer training target and handles the non-texture regions. Our experimental evaluation demonstrates that the result of our method is comparable to fully supervised methods on the NYU Depth V2 benchmark. To the best of our knowledge, this is the first quantitative result of purely unsupervised learning method reported on indoor datasets.

1. Introduction

Reconstructing the structures of scenes from RGB images has long been a popular research topic. Depth estimation is an important step to reconstruct the scenes. It is easy for human to perceive the depth of scenes from a RGB image since we have prior knowledge about the sizes of common objects and their layouts, but it is difficult for computer to estimate accurate depth map from a single image due to scale ambiguity. Classical methods like Structure-from-Motion and Stereo Matching [14, 18] were proposed and achieved plausible results. In recent years, due to the surge of deep learning, convolutional neural network was introduced to predict depth from monocular RGB images for its

powerful capacity of feature extraction, which leads to great progress in this field. These methods [31, 11, 10], treat the neural networks as blackbox with strong fitting ability and use the collected ground-truth annotation to supervise the training. However, these fully supervised methods are limited by the huge demand of training samples.

Recently unsupervised depth learning [48, 17, 47] has been proposed and attracted more and more interests. This method is similar to traditional Structure-from-Motion that leverages the disparity information contained in videos to supervise the networks' training. The key idea is novel view synthesis based on simultaneously estimated depth of scene and ego-motion of camera. The appearance difference between the synthesized view and real view is used as supervisory signal for the entire training pipeline. These unsupervised methods do not need ground-truth annotation and achieve remarkable results on driving scenes like KITTI and Cityscapes.

However, there still exist great challenges when directly applying this technology in indoor environments. In our experiments, we observed that the same model and same training setting, which is able to achieve state-of-the-art performance on KITTI, soon collapse when training on indoor datasets NYU V2[39] and Scannet [8]. The reason is that indoor environments are more complicated than urban driving scenes. The main problems can be summarized below:

1) Large areas of non-texture regions. Unlike fully supervised methods in which each pixel has ground-truth supervision, the supervisory signal of unsupervised learning only comes from the appearance difference between images themselves. Non-texture regions seriously hinder the training since in these regions the photometric loss is always close to zero. However we observed that there are considerable amount of images which have more than 50% non-texture areas in indoor datasets. White wall and carpets are

*Work done as an intern at MSRA.

fairly common non-texture objects.

2) More complex ego-motion of handheld camera. Generally indoor datasets are collected by handheld cameras, which means that the ego-motion of consecutive frames is more complex than driving scenes where the cars are mostly just moving forward. Especially, we can not infer the depth of scene theoretically from sequences of pure rotation. The existence of large amount of training samples with pure rotation will overwhelm the entire training process.

These reveal that the existing training pipeline needs to be revamped to be applied in more general scenes. In this paper, we propose a new optical-flow based training paradigm, which focuses on the most important part, i.e., the supervisory signal of unsupervised depth learning. This new pipeline uses the optical flow results generated by a flow estimation network as supervision and it is easier for the training to converge. The key component is a specially designed network which is responsible for estimating the optical flow between consecutive frames in a sparse-to-dense propagation manner. This unsupervised flow network is able to handle the non-texture regions and produce plausible optical flow results. Then this network can be used as a teacher simultaneously training the DepthNet and the PoseNet. We also improve the existing PoseNet and make it easier to learn the complex ego-motion of handheld camera.

Our evaluation on NYU Depth V2 benchmark demonstrates that the result of our method is comparable to fully supervised methods.

2. Related Work

Supervised Depth Estimation Estimating the depth from a single image has been long studied. Recently due to the success of deep learning [27, 28], many networks for depth estimation have been proposed [11, 10, 29, 31, 13, 24]. Eigen et al. [11] applied multi-scale networks which first estimate coarse depth by a coarse-scale network and refine it by another network. CRF is also introduced to this task and used as a postprocessing module in a model [45]. Fu et al. [13] treated the depth estimation as a classification problem instead of a regression problem. All these methods have powerful capacity and achieve very good performance on both indoor datasets like NYU V2 [39], Scannet [8] and outdoor datasets like KITTI [15], Make3D [37, 38]. However these methods rely on large-scale datasets with depth labels.

Unsupervised Depth Learning To get rid of the need of ground-truth depth annotation, unsupervised depth learning methods have been proposed. These methods leverage either stereo images [17, 46] or videos [48, 47] as training data. Godard et al. [17] first proposed to use the left-right consistency of stereo images to train a depth estimation network. Zhou et al. [48] applied two networks which jointly

estimate the depth and ego-motion of camera to learn the depth from videos. Wang et al. [42] discarded the pose network and directly computed the pose by visual odometry method. Casser et al. [3] leveraged the additional instance segmentation masks to model the dynamic objects. These methods have achieved tremendous success in outdoor scenes like KITTI [15] and CityScapes [7]. However, only some of these works demonstrated the sample predicted results of indoor scenes using the networks trained on KITTI, and there is no quantitative result reported on typical indoor datasets. In our experiments, we also faced great challenges when directly using the previous methods to train on indoor scenes.

Unsupervised Optical Flow Learning Based on the same photometric supervisory signal as unsupervised depth learning, unsupervised optical flow learning methods have also been proposed. Yu et al. [21] and Ren et al. [34] proposed FlowNet-based [19] architecture for unsupervised optical flow learning. Meister et al. [33] proposed a Bidirectional Census Loss to handle occlusion/disocclusion. Although these methods perform well on synthetic datasets, non-texture regions, dynamic objects and occlusion are still intractable problems. Our new flow network is different from previous architecture, which leverages the sparse flow seeds generated by traditional feature matching methods and progressively propagates them to the entire image. Non-texture regions are handled well by this mean.

3. Approach

3.1. Overall Pipeline

As shown in Figure 1 (a), previous unsupervised depth learning pipeline generally consists of two modules: DepthNet and PoseNet. During training, both networks simultaneously estimate the depth of the scene and the ego-motion of the camera. Once the depth map of RGB image is estimated, we can backproject the pixels on image plane to 3D coordinates with known camera intrinsic. Then with the estimated motion the 3D point cloud can be transformed to another view. This view transformation can be formulated as below:

$$p_{s \rightarrow t} = K T_{t \rightarrow s} D_t(p_t) K^{-1} p_t \quad (1)$$

where K denotes the camera intrinsic, $T_{t \rightarrow s}$ denotes the estimated transformation matrix from view t to view s , D_t denotes the estimated depth, p_t and $p_{s \rightarrow t}$ denote the homogeneous coordinates of a pixel in view t and view s respectively. Then we can obtain the 2D rigid flow from view t to view s :

$$f_{t \rightarrow s}(p_t) = p_{s \rightarrow t} - p_t \quad (2)$$

Once the rigid flow between two views is estimated, we can synthesize the image $I_{s \rightarrow t}$ by differentiable inverse warp-

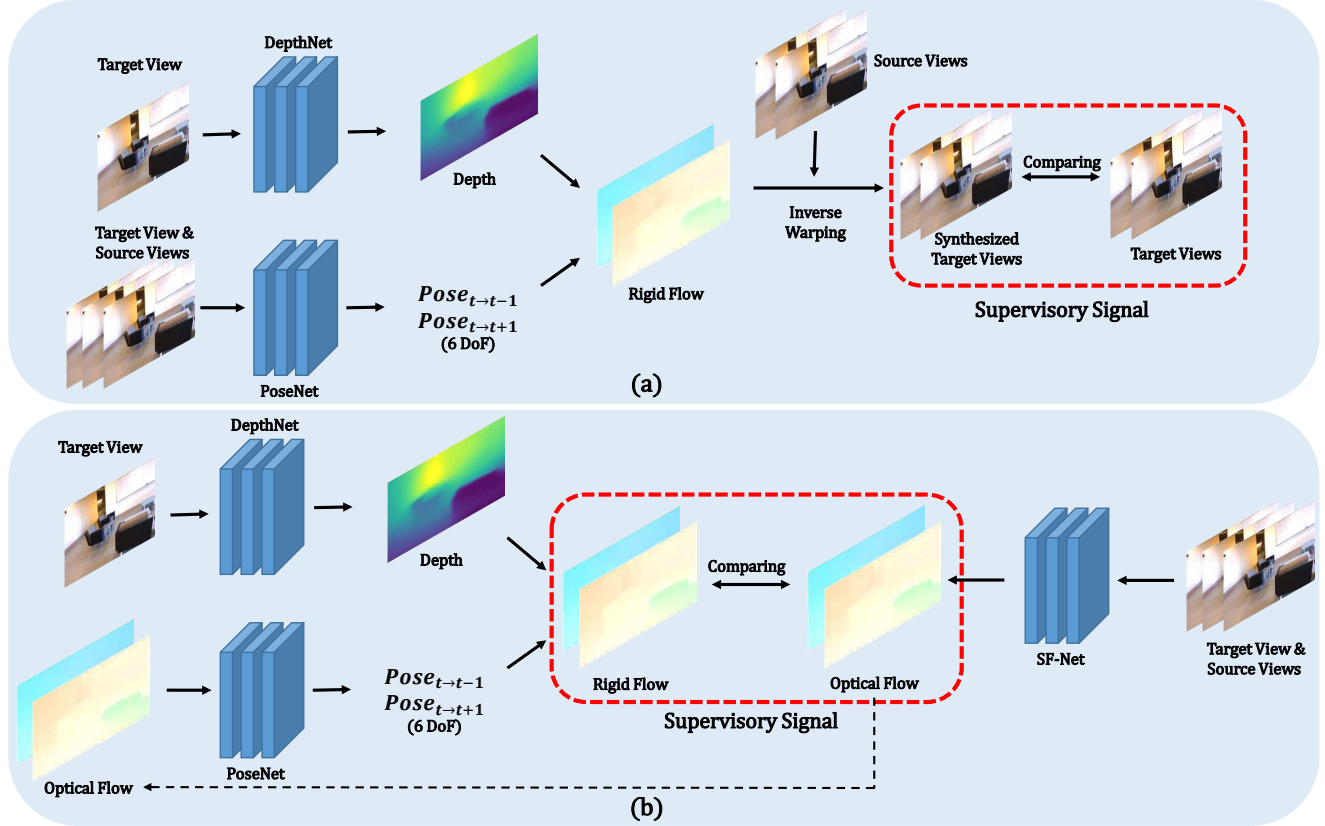


Figure 1. Overview of previous pipeline (a) and our pipeline (b). The supervisory signal of previous pipeline is based on appearance matching of two images, which is unstable and suffers from non-texture regions. We use the optical flow results generated by our flow estimation network SF-Net as supervision and modify the input of PoseNet.

ing [20] from source view. The appearance difference between the synthesized image $I_{s \rightarrow t}$ and the real image I_t is used as supervisory signal for the entire pipeline.

As shown in the red frame in Figure 1 (a), the photometric loss computed based on the synthesized image and the real image supervises the training of both networks. This is equivalent to finding the best match point in source view for each pixel in target view, which is similar to Stereo Matching. Once the rigid flow is predicted perfectly, the synthesized image matches the real image completely (if there is no occlusion and dynamic object). However, this appearance-based supervisory signal is indirect and susceptible to non-texture regions. More specifically, for each pixel there is no explicit target position that it should match in another view. The optimization target of this proxy objective is just to minimize the appearance difference between two images. This paradigm is difficult to work in indoor environments since there is a lot of non-texture stuff in our daily scenes, where the appearance difference is always close to zero, which can not provide valid and strong signal for the training.

To overcome the problem, we first look back to the entire pipeline. It can be separated into two stages: 1) composing

the rigid flow based on jointly estimated depth and pose. 2) using the rigid flow to synthesize novel image and computing the loss (corresponding to the left and right parts in Figure 1 (a)). The objective function can be briefly written as:

$$L = |I_t - I_{s \rightarrow t}| \quad (3)$$

where I_t denotes target view image and $I_{s \rightarrow t}$ denotes synthesized image.

Our main contribution is that instead of using this indirect proxy supervision, we provide an explicit optical flow target to supervise the estimated rigid flow as shown in Figure 1 (b). The optical flow target is obtained by a sparse-to-dense flow estimation network (named hereafter as SF-Net), which will be elaborated on in the next subsection. Therefore the objective function is modified as:

$$L = |f_{t \rightarrow s}(p_t) - f'_{t \rightarrow s}(p_t)| \quad (4)$$

where $f'_{t \rightarrow s}(p_t)$ denotes the optical flow result from SF-Net. This modification is significant since it transforms the unsupervised learning to “fully supervised” learning and reduces the difficulty of training. The words “fully su-

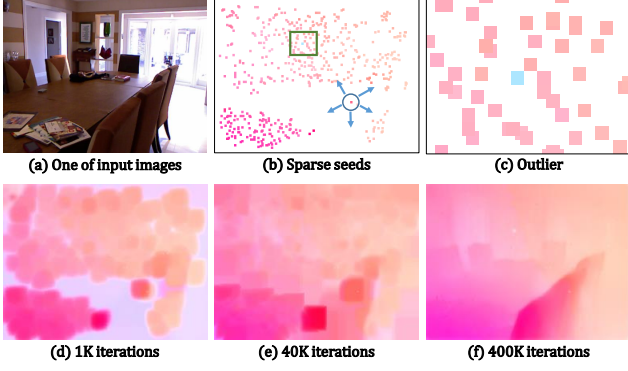


Figure 2. Illustration of our SF-Net. (a) One of input images. (b) Sparse seeds generated by SURF [1]. The sizes of sparse points have been enlarged for better visualization. (c) The green frame in (b), which represents an outlier (blue point) in sparse seeds. (d)(e)(f) Visualization of a training sample at different stages. Our SF-Net works in a sparse-to-dense propagation manner which progressively propagates the sparse flow from textured regions to non-texture regions. The negative effects induced by outliers are also suppressed during training.

pervised” here only mean more explicit supervisory signal while this signal is still obtained unsupervisedly. Photometric loss is not able to penalize the incorrect prediction in non-texture regions, i.e., the supervisory signal fails in those areas (common in indoor scenes). By contrast, the optical flow generated by SF-Net provides unique target for each pixel, which is a very strong supervisory signal. Our experiments show that this new pipeline is able to handle the indoor environments with large areas of non-texture regions and achieve plausible result.

3.2. SF-Net

The core of this new training paradigm is the part of SF-Net. Before the introduction of the SF-Net, we first explain the principle of unsupervised optical flow learning. The key idea is novel view synthesis, which is the same as unsupervised depth learning. The network predicts a flow map for image I_t in target view and use this dense flow to synthesize the image $I_{s \rightarrow t}$ by differentiable inverse warping [20] from source view. Then the appearance difference between the synthesized image $I_{s \rightarrow t}$ and the real target image I_t is used as supervisory signal for the training.

However, since both unsupervised depth learning and optical flow learning leverage the same proxy supervisory signal, they also suffer from the same problems. Previous methods usually impose smoothness constraint on the flow prediction. The reason is that it is easy for the network to produce correct result in textured region like corner and border, we expect this correct prediction to guide its neighborhood where the prediction is incorrect. But in indoor environments where non-texture regions cover large areas, the correct prediction may be overwhelmed by the incorrect

predictions on the contrary.

In order to overcome this problem, we propose an active propagation approach. Instead of only using the weak smoothness constraint, we actively propagate the sparse initial seeds at textured regions to the entire image as shown in Figure 2. The key idea is that we do not need to generate the dense flow maps from scratch since we can leverage the traditional feature matching algorithms like SURF [1] to generate sparse corresponding points. The displacement of the corresponding points (Figure 2 (b)) are considered as initial flow seeds and propagated to the entire region. In addition, the network is able to suppress the mismatches existing in the corresponding points during training. In this way, plausible results can also be generated in the interior of non-texture regions.

As for the method of propagation, we adopt the architecture of CSPN proposed by Cheng et al.[5, 6] for its efficiency. This method diffuses the information of the center pixel to its eight neighborhood iteratively in the form of convolution. Our adapted network is a very simple and common encoder-decoder architecture, which takes stacked RGB images and sparse seeds as input and outputs two results: one is the coarse optical flow F_0 and another is the transformation kernels $\hat{K}_{i,j}$ with $k^2 - 1$ channels, where k denotes the kernel size. Then the coarse optical flow is refined iteratively as in [5, 6]:

$$F_{i,j,t+1} = \sum_{a,b=-(k-1)/2}^{(k-1)/2} K_{i,j}(a,b) \odot F_{i-a,j-b,t} \quad (5)$$

where \odot denotes element-wise product and:

$$K_{i,j}(a,b) = \frac{\hat{K}_{i,j}(a,b)}{\sum_{a,b \neq 0} |\hat{K}_{i,j}(a,b)|} \quad (6)$$

$$K_{i,j}(0,0) = 1 - \sum_{a,b \neq 0} K_{i,j}(a,b) \quad (7)$$

Before each propagating operation, we fix the sparse seeds in order to guarantee that our propagated flows have the exact same value at those valid pixels in the sparse flow map:

$$F_{i,j,t+1} = (1 - m_{i,j})F_{i,j,t+1} + m_{i,j}F_{i,j}^s \quad (8)$$

where $F_{i,j}^s$ denotes the sparse flow with empty positions filled with zero, $m_{i,j}$ is an indicator for the availability of sparse flow at (i,j) . In our training setting, the kernel size k is set to 3 and the max steps of iterations are 16.

3.3. PoseNet

The PoseNet is also an important component of the unsupervised depth learning pipeline, which is an application of deep learning in visual odometry and responsible for estimating the pose in 6 degrees of freedom (DoF) between

two images. In driving scenes like KITTI the poses are fairly simple and the cars in most images are just driving forward. But in indoor environments, the images are typically collected by handheld cameras, which means more complex ego-motion and raises the difficulty of learning for PoseNet.

There are also some approaches proposed to get rid of the PoseNet. For instance, Mahjourian et al. [32] used Iterative Closest Point (ICP) [2, 4, 36] to compute a transformation that minimizes point-to-point distances between corresponding points. Wang et al. [42] used direct visual odometry (DVO) [40] to obtain camera pose from predicted depth and images. But in our experiments these network-free methods collapsed during training on indoor datasets since they rely on good initial depth with low level of noise. Collapse means all the predictions of depth converge to a constant value. The estimation of poses can also be considered as a Perspective-n-Point (PNP) problem because we have predicted depth and dense matching simultaneously. We tried to implement the EPNP [25] algorithm on GPU to directly compute camera pose from predicted depth and flow in real time but the training collapsed again. These trials indicate that the use of PoseNet is necessary since the prediction of PoseNet is based on the entire dataset statistically and it is not likely to be dominated by individual training samples.

Therefore we rethink the working principle of the PoseNet. How does this blackbox estimate the pose from stacked RGB images? A reasonable speculation is that it first finds corresponding points in two images internally and infers the pose from the displacements of the corresponding points according to an unknown rule. However, there is no need for the PoseNet to match the pixels again since we already have the optical flow result, i.e., the dense matching. So we propose to use the flow result produced by SF-Net as the input of PoseNet instead of RGB images.

This modification is equivalent to separate the unsupervised pose estimation into two stages: estimating the optical flow between two frames first and then inferring the pose based on the flow. It is unnecessary to do this separation if there exist ground-truth labels. But it enhances the interpretability of the PoseNet and reduces the difficulty of unsupervised learning. The significant improvement in the experimental results reported in Section 4.2 also supports our speculation.

3.4. Pure Rotation

There still exists a problem that should not be ignored: pure rotation. This problem does not exist in driving scenes since previous methods generally remove static frames during preprocessing but is common in indoor datasets. Theoretically image pairs with pure rotation do not contain depth information and are harmful to the training. It is crucial to

first filter out the training samples with pure rotation. It can be derived that the relationship between corresponding points with pure rotation can be fitted by a homography matrix H :

$$H = K R K^{-1} \quad (9)$$

where K denotes the camera intrinsic and R denotes the rotation matrix.

So for each image pair, we use the dense optical flow generated by SF-Net to compute its homography matrix with RANSAC [12]. If the ratio of outliers is lower than a pre-set threshold (20% in our setting, which means more than 80% pixels can be fitted by a homography matrix), we consider the pose of this image pair as pure rotation and discard it. After this filtering process, about 30% of the images in NYU V2 are discarded.

3.5. Loss Function

3.5.1 Loss Function of SF-Net

For the training of SF-Net, we use photometric loss and smoothness loss.

Photometric Loss This loss function computes the appearance difference between two images. We adopt the same setting as [17] which combines L1 loss and the Structural Similarity (SSIM) [44]. Besides per-pixel minimum trick proposed by [16] is also adopted which is aimed to handle occlusion/disocclusion. It can be written as:

$$L_{ph} = \sum_p \min_s (\alpha SSIM(I_t, I_{s \rightarrow t}) + (1 - \alpha) |I_t(p) - I_{s \rightarrow t}(p)|) \quad (10)$$

where p indexes over pixel coordinates, s denotes the index of source views, α is set to 0.5.

Smoothness Loss Although SF-Net adopts a sparse-to-dense training scheme, we also use edge-aware flow smoothness loss to suppress the mismatches in initial seeds:

$$L_{smooth} = \sum_p |\nabla F(p)| \cdot \left(e^{-|\nabla I(p)|} \right)^T \quad (11)$$

where ∇ is the vector differential operator, and T denotes the transpose of image gradient weighting. So the total loss of SF-Net is:

$$L = \lambda_1 L_{ph} + \lambda_2 L_{smooth} \quad (12)$$

3.5.2 Loss Function of DepthNet and PoseNet

For the training of DepthNet and PoseNet, the loss function consists of four terms.

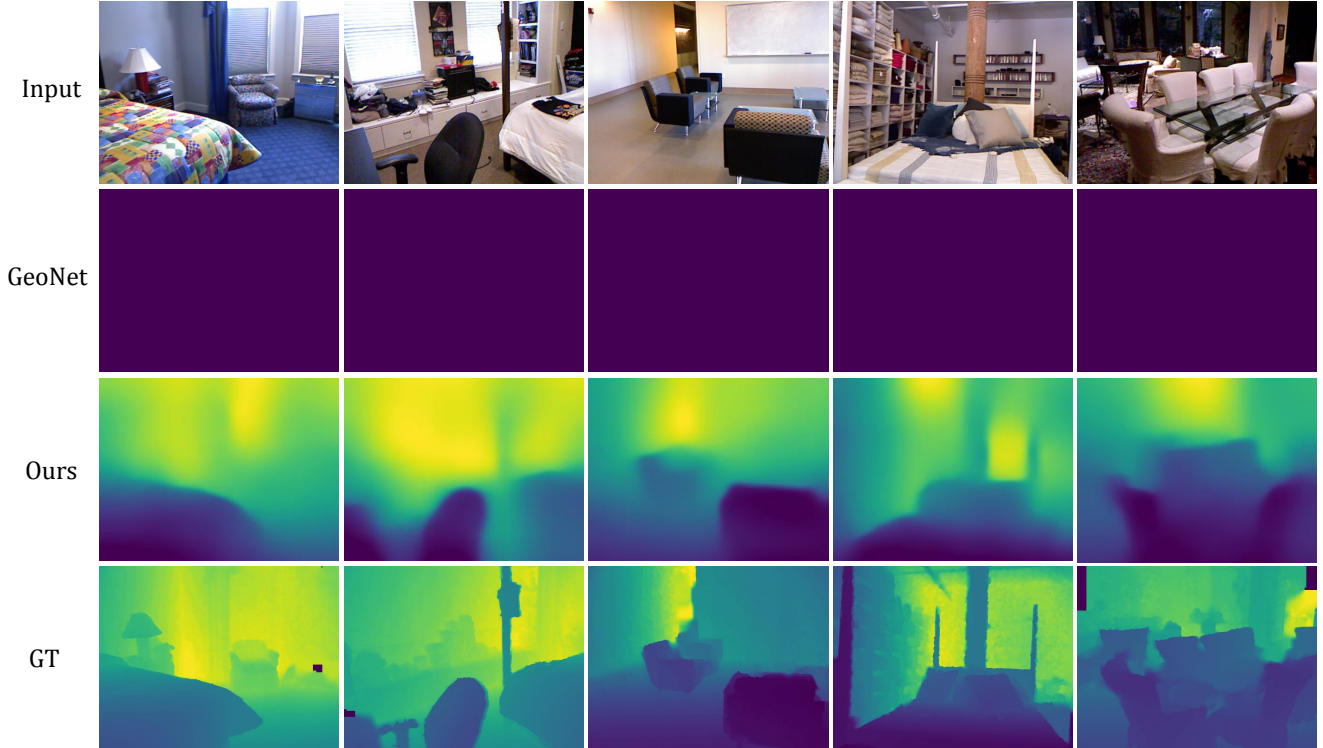


Figure 3. Qualitative comparison between GeoNet [47], ours and ground-truth depth. We directly use the original code of GeoNet, which is able to achieve state-of-the-art performance on KITTI, but collapsed during training on NYU V2 [39].

Method	Supervision	Accuracy metric			Error metric		
		$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$	rel	log10	rms
Make3D [38]	✓	0.447	0.745	0.897	0.349	-	1.214
Depth Transfer [22]	✓	-	-	-	0.35	0.131	1.2
Liu et al. [31]	✓	-	-	-	0.335	0.127	1.06
Ladicky et al. [23]	✓	0.542	0.829	0.941	-	-	-
Li et al. [26]	✓	0.621	0.886	0.968	0.232	0.094	0.821
Wang et al. [43]	✓	0.605	0.890	0.970	0.220	-	0.824
Roy et al. [35]	✓	-	-	-	0.187	-	0.744
Liu et al. [30]	✓	0.650	0.906	0.976	0.213	0.087	0.759
Li et al. [29]	✓	0.788	0.958	0.991	0.143	0.063	0.635
MS-CRF [45]	✓	0.811	0.954	0.987	0.121	0.052	0.586
DORN [13]	✓	0.828	0.965	0.992	0.115	0.051	0.509
Ours (baseline†)	×	0.511	0.779	0.904	0.331	0.127	1.000
Ours	×	0.674	0.900	0.968	0.208	0.086	0.712

Table 1. Comparison to existing methods on NYU V2 [39]. † denotes the model that collapsed during training. Since we adopt the scale normalization proposed by [42] to avoid the shrinking of depth, all the predictions are normalized to 1 meter when the model collapses. All the other methods in the table are fully supervised by depth annotations.

Rigid Flow Loss This term directly uses the optical flow result produced by SF-Net as supervision of the synthesized rigid flow. We use berHu [24] norm $\|\cdot\|_\delta$ to measure the deviation:

$$L_{flow} = \|f_{t \rightarrow s}(p_t) - f'_{t \rightarrow s}(p_t)\|_\delta \quad (13)$$

The other two terms L_{s_depth} and L_{s_normal} are smoothness constraints and are similar to Equation 11. But they are imposed on the predicted depth and corresponding normal

Method	Supervision	Filtering of Pure Rotation	Accuracy metric			Error metric		
			d1	d2	d3	rel	log10	rms
DepthNet + R-PoseNet †	RGB	×	0.511	0.779	0.904	0.331	0.127	1.000
DepthNet + R-PoseNet †	RGB	✓	0.511	0.779	0.904	0.331	0.127	1.000
DepthNet + F-PoseNet †	Flow (w/o propagation)	×	0.511	0.779	0.904	0.331	0.127	1.000
DepthNet + F-PoseNet	Flow (w/o propagation)	✓	0.596	0.862	0.951	0.257	0.102	0.841
DepthNet + R-PoseNet	Flow (w/ propagation)	✓	0.578	0.836	0.938	0.273	0.108	0.910
DepthNet + F-PoseNet	Flow (w/ propagation)	✓	0.674	0.900	0.968	0.208	0.086	0.712

Table 2. Evaluation of each component on NYU V2’s test split. † denotes the model that collapsed during training. R-PoseNet denotes the PoseNet with RGB images input and F-PoseNet denotes the PoseNet with flow input. $\delta < 1.25$, $\delta < 1.25^2$, $\delta < 1.25^3$ are abbreviated to d1, d2, d3 due to space limitation.

Method	Dataset	Accuracy		Error		
		d1	rel	sq rel	rms	
Ours(baseline)†	S	0.631	0.238	0.190	0.570	
Ours(F-PoseNet)†	S	0.631	0.238	0.190	0.570	
Ours(F-Sup+R-PoseNet)	S	0.682	0.206	0.134	0.491	
Ours(F-Sup+F-PoseNet)	S	0.710	0.190	0.124	0.465	
Ours(baseline)	K	0.833	0.132	1.016	5.506	
Ours(F-Sup)	K	0.837	0.136	1.110	5.327	
Ours(F-PoseNet)	K	0.836	0.130	1.001	5.294	

Table 3. Evaluation of depth on Scannet and KITTI datasets. K:KITTI, S:Scannet, †:collapse, F-sup:flow supervision.

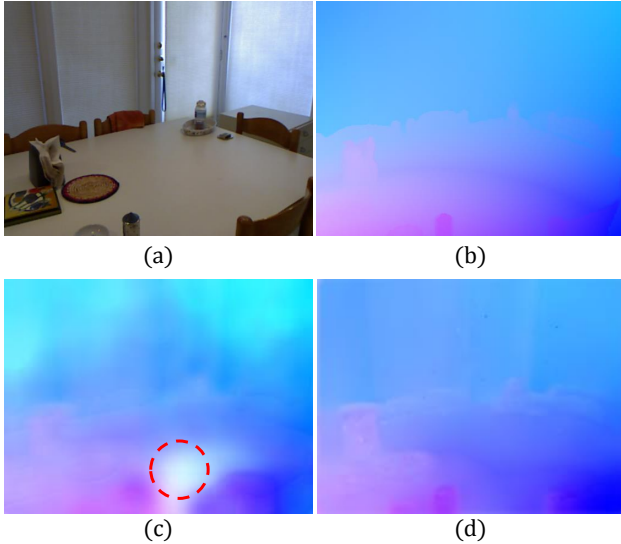


Figure 4. (a) A typical image with large areas of non-texture regions. (b) GT flow. (c) Flow without propagation. (d) Flow with propagation. Previous unsupervised flow learning methods only impose smoothness constraint on non-texture regions, which are prone to consider these areas as static (red circle in (c)).

which is computed from the depth. The last term L_{ph} is the same as Equation 10. So the total loss of DepthNet and PoseNet is

$$L = \lambda_1 L_{flow} + \lambda_2 L_{ph} + \lambda_3 L_{s_depth} + \lambda_4 L_{s_normal} \quad (14)$$

Method	w/o Propagation	w/ Propagation	PWC-Net
EPE	7.409	3.602	3.279

Table 4. Average endpoint error (EPE) flow results on NYU validation split. The result of PWC-Net is directly generated by PWC-Net [41] pretrained on FlyingChairs [9] with supervision but without finetuning on NYU V2. The result obtained by SF-Net with propagation is close to supervised result obtained by PWC-Net.

4. Experiments

4.1. Comparisons on NYU V2 Dataset

The NYU Depth v2 [39] dataset contains 582 indoor video scenes taken with a Microsoft Kinect camera and the training split contains 283 scenes (about 230K images). To train the DepthNet and PoseNet, we first use the method mentioned in Section 3.4 to filter out the image pairs with pure rotation. About 30% images are discarded and finally we use about 180K images for training. We fix the length of training image sequences to be 3 frames for all the three networks, and treat the central frame as the target view and the ± 10 frames as the source views. We only use raw RGB image sequences for training and the images are resized to 192×256 . Visualization is shown in Figure 3 and quantitative evaluation is reported in Table 1. The baseline model only consists of DepthNet and PoseNet with RGB images input and does not use flow as supervision.

4.2. Ablation Study

In this subsection, we individually evaluate the effects of the four components in our pipeline: 1) The flow supervision. 2) The filtering of pure rotation. 3) The propagation design for the SF-Net. 4) The input of the PoseNet.

As shown in Table 2, the results of rows 3,6 and rows 4,5 indicate that the flow supervision and filtering of pure rotation are indispensable in order to get rid of the collapse phenomenon. It also shows that both the F-PoseNet and the propagation of SF-Net improve the performance significantly. The result without propagation means the SF-Net takes RGB images as input and only output optical flow,

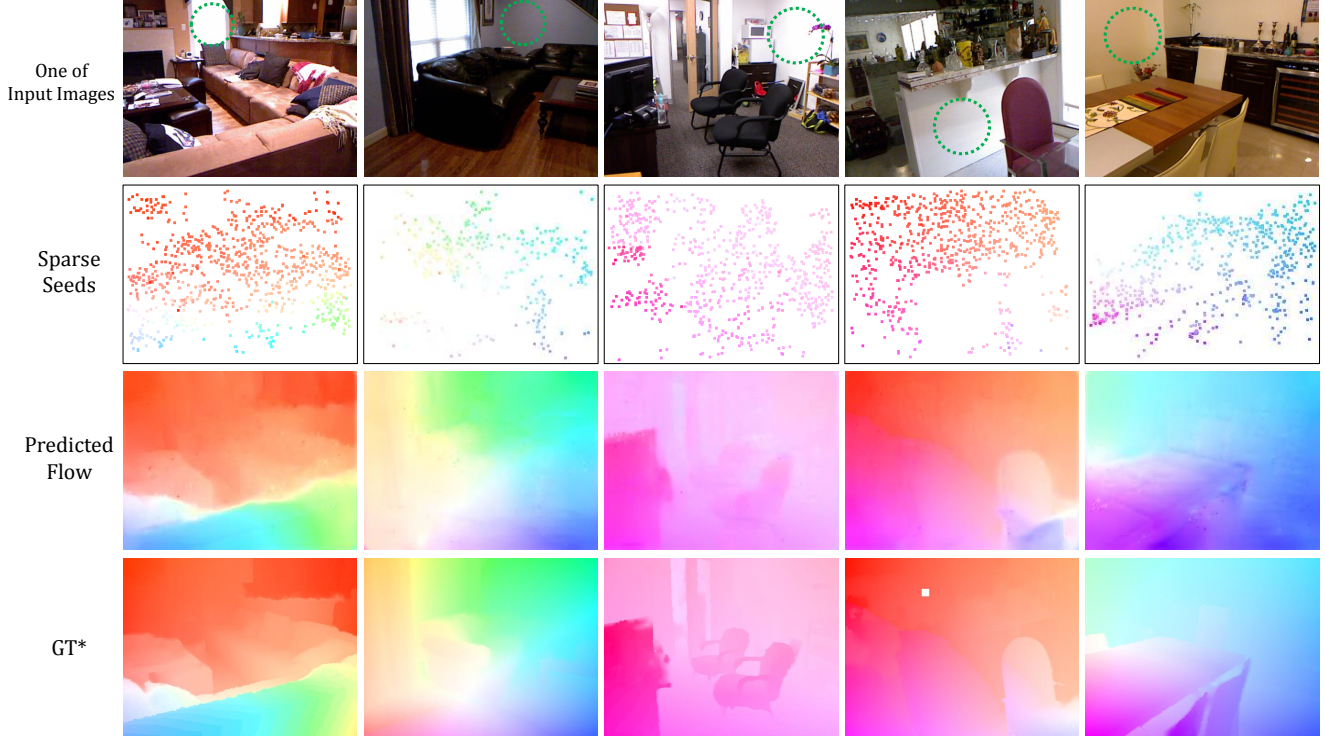


Figure 5. Visualization of the results generated by SF-Net on NYU V2. NYU depth dataset does not contain optical flow annotations. * means that the flow results are computed with ground-truth depth by the method mentioned in Section 4.2. Our SF-Net handles the non-texture regions (green circles) well. The sizes of sparse points have been enlarged for better visualization.

and the other parts in the network remains the same.

NYU depth dataset does not contain optical flow annotations. But we can compute the rigid flow from the ground-truth depth. We first use the EPNP [25] algorithm to solve the pose between two images with depth annotations and sparse matchings. The sparse matchings are also obtained by SURF [1]. Then the estimated pose and ground-truth depth can be used to compose the rigid flow. We consider these estimated rigid flows as ground-truth flow labels.

In order to quantitatively evaluate the performance of SF-Net, we randomly choose 1000 image pairs as validation set and exclude them during training. The quantitative result is shown in Table 4 and qualitative comparison is shown in Figure 4. With the help of the propagation of sparse seeds, SF-Net handles the non-texture regions well and achieves reasonably good results as shown in Figure 5.

4.3. Evaluations on Scannet and KITTI

As shown in Table 3, we also evaluate our method on Scannet and KITTI datasets. Since Scannet is a large dataset, we train and test our model on a subset of Scannet for efficiency, which contains 40 scenes and about 70K images. The flow supervision is proposed to address the non texture problems in indoor scenarios, but has its own limitation. When applied on KITTI, the accuracy of the flow becomes the bottleneck. However, F-PoseNet still has gain.

5. Conclusion

In this paper, we propose a new unsupervised depth learning framework which reduces the difficulty for the network to learn and is able to work indoor. We also propose a sparse-to-dense unsupervised flow estimation network that addresses the intractable non-texture area problem. More importantly, the results of our approach demonstrate that the technology of unsupervised depth learning is not only able to work in driving scenes but also has the capacity to be applied in more general scenes. This is an important step towards exploring the innumerable videos available on the Internet for the training of deep learning.

Limitation Although our approach is able to handle the non-texture regions in most cases, it relies on the propagation of sparse seeds. In some special situations where non-texture areas are extremely large and the identified corresponding key points are very sparse, it is difficult for our model to predict correct results. In addition, the generated flows usually have blurry boundary. When these blurry flows are used to supervise the training of DepthNet, the obtained depth maps have blurrier boundary. This also limits the performance of our model. We will address these problems in the future work.

References

- [1] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *European conference on computer vision*, pages 404–417. Springer, 2006.
- [2] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures*, volume 1611, pages 586–607. International Society for Optics and Photonics, 1992.
- [3] Vincent Casser, Soeren Pirk, Reza Mahjourian, and Anelia Angelova. Unsupervised learning of depth and ego-motion: A structured approach. In *Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*, 2019.
- [4] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Image and vision computing*, 10(3):145–155, 1992.
- [5] Xinjing Cheng, Peng Wang, and Ruigang Yang. Depth estimation via affinity learned with convolutional spatial propagation network. In *European Conference on Computer Vision*, pages 108–125. Springer, Cham, 2018.
- [6] Xinjing Cheng, Peng Wang, and Ruigang Yang. Learning depth with convolutional spatial propagation network. 2018.
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, pages 3213–3223, 2016.
- [8] 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.
- [9] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. v.d. Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [10] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, 2015.
- [11] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NIPS*, 2014.
- [12] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [13] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *CVPR*, 2018.
- [14] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015.
- [15] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012.
- [16] Clément Godard, Oisín Mac Aodha, and Gabriel Brostow. Digging into self-supervised monocular depth estimation. *arXiv preprint arXiv:1806.01260*, 2018.
- [17] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017.
- [18] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [19] 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.
- [20] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *NIPS*, 2015.
- [21] J Yu Jason, Adam W Harley, and Konstantinos G Derpanis. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *European Conference on Computer Vision*, pages 3–10. Springer, 2016.
- [22] Kevin Karsch, Ce Liu, and Sing Bing Kang. Depth transfer: Depth extraction from video using non-parametric sampling. *PAMI*, 2014.
- [23] Lubor Ladicky, Jianbo Shi, and Marc Pollefeys. Pulling things out of perspective. In *CVPR*, 2014.
- [24] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *3D Vision (3DV)*, 2016.
- [25] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnp: An accurate o(n) solution to the pnp problem. *International journal of computer vision*, 81(2):155, 2009.
- [26] Bo Li, Chunhua Shen, Yuchao Dai, Anton Van Den Hengel, and Mingyi He. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1119–1127, 2015.
- [27] Hongyang Li, Bo Dai, Shaoshuai Shi, Wanli Ouyang, and Xiaogang Wang. Feature Intertwiner for Object Detection. In *ICLR*, 2019.
- [28] Hongyang Li, David Eigen, Samuel Dodge, Matt Zeiler, and Xiaogang Wang. Finding Task-relevant Features for Few-shot Learning by Category Traversal. In *CVPR*, 2019.
- [29] Jun Li, Reinhard Klein, and Angela Yao. A two-streamed network for estimating fine-scaled depth maps from single rgb images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3372–3380, 2017.
- [30] 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, 2016.
- [31] Miaomiao Liu, Mathieu Salzmann, and Xuming He. Discrete-continuous depth estimation from a single image. In *CVPR*, 2014.
- [32] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In *CVPR*, 2018.
- [33] Simon Meister, Junhwa Hur, and Stefan Roth. Unflow: Unsupervised learning of optical flow with a bidirectional cen-

- sus loss. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [34] Zhe Ren, Junchi Yan, Bingbing Ni, Bin Liu, Xiaokang Yang, and Hongyuan Zha. Unsupervised deep learning for optical flow estimation. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
 - [35] Anirban Roy and Sinisa Todorovic. Monocular depth estimation using neural regression forest. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5506–5514, 2016.
 - [36] Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the icp algorithm. In *3dim*, page 145. IEEE, 2001.
 - [37] Ashutosh Saxena, Sung H Chung, and Andrew Y Ng. Learning depth from single monocular images. In *NIPS*, 2006.
 - [38] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Make3d: Learning 3d scene structure from a single still image. *PAMI*, 2009.
 - [39] 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.
 - [40] Frank Steinbrücker, Jürgen Sturm, and Daniel Cremers. Real-time visual odometry from dense rgb-d images. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 719–722. IEEE, 2011.
 - [41] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. 2018.
 - [42] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *CVPR*, 2018.
 - [43] Peng Wang, Xiaohui Shen, Zhe Lin, Scott Cohen, Brian Price, and Alan L Yuille. Towards unified depth and semantic prediction from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2800–2809, 2015.
 - [44] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 2004.
 - [45] Dan Xu, Elisa Ricci, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Multi-scale continuous crfs as sequential deep networks for monocular depth estimation. In *CVPR*, 2017.
 - [46] Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, and Ram Nevatia. Every pixel counts: Unsupervised geometry learning with holistic 3d motion understanding. In *ECCV*, 2018.
 - [47] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *CVPR*, 2018.
 - [48] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017.