

# Self-Supervised Learning of Depth and Motion Under Photometric Inconsistency

Tianwei Shen  
HKUST

tshenaa@cse.ust.hk

Shiwei Li  
HKUST

slibc@cse.ust.hk

Lei Zhou  
HKUST

lzhouai@cse.ust.hk

Jiahui Zhang  
Tsinghua University

zjhthu@gmail.com

Zixin Luo  
HKUST

zluoag@cse.ust.hk

Tian Fang  
Altizure

fangtian@altizure.com

Yao Yao  
HKUST

yyaoag@cse.ust.hk

Long Quan  
HKUST

quan@cse.ust.hk

## Abstract

The self-supervised learning of depth and pose from monocular sequences provides an attractive solution by using the photometric consistency of nearby frames as it depends much less on the ground-truth data. In this paper, we address the issue when previous assumptions of the self-supervised approaches are violated due to the dynamic nature of real-world scenes. Different from handling the noise as uncertainty, our key idea is to incorporate more robust geometric quantities and enforce internal consistency in the temporal image sequence. As demonstrated on commonly used benchmark datasets, the proposed method substantially improves the state-of-the-art methods on both depth and relative pose estimation for monocular image sequences, without adding inference overhead.

## 1. Introduction

The joint learning of depth and relative pose from monocular videos [47, 51, 54] has been an active research area due to its key role in *simultaneous localization and mapping* (SLAM) and *visual odometry* (VO) applications. The simplicity and the unsupervised nature make itself a potential replacement for traditional approaches that involve complicated geometric computations. Given adjacent frames, this approach uses convolutional neural networks (CNNs) to jointly predict the depth map of the target image and the relative poses from the target image to its visible neighboring frames. With the predicted depth and relative poses, photometric error is minimized between the original target image and the synthesized images formed by bilinear-sampling [19] the adjacent views.

However, several existing problems hinder the performance of this approach. First, the photometric loss requires the modeling scene to be static without non-Lambertian surfaces or occlusions. This assumption is often violated in

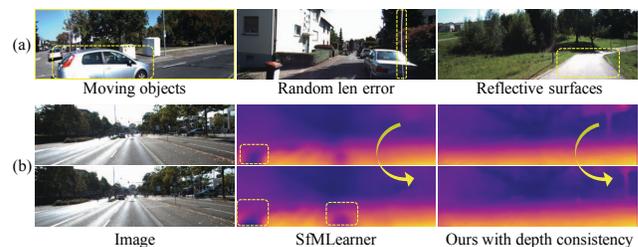


Figure 1. (a) Violations of the photometric consistency in KITTI [13]. (b) The lack of depth consistency leads to erroneous ‘black holes’ (middle [54]) in the depth estimation.

street-view datasets [6, 13] with moving cars and pedestrians (see Figure 1(a) for some failure cases). To this end, we need other stable supervisions that are less affected when the photometric consistency is invalid. Second, as the monocular depth inference considers only single images, there is no guarantee that adjacent frames would have consistent depth estimation. This increases the chance that the inferred outcome would contain noisy depth values, and ignores the information from adjacent views when it is readily available. In addition, using pure color information is subject to the well-known gradient locality issue [5]. When image regions with vastly different depth ranges have the similar appearance (e.g. the road in Figure 1(b)), gradients inferred from photometric information are not able to effectively guide the optimization, leading to erroneous patterns such as ‘black holes’ (erratic depth).

In this paper, we propose a novel formulation that emphasizes various consistency constraints of deep interplay between depth and pose, seeking to resolve the photometric inconsistency issue. We propose the geometric consistency from sparse feature matches, which is robust to illumination changes and calibration errors. We also show that enforcing the depth consistency across adjacent frames significantly improves the depth estimation with much fewer noisy pix-

els. The geometric information is implicitly embedded into neural networks and does not bring overhead for inference.

The consistency of multi-view geometry has been widely applied to and even forms the basis for many sub-steps in SfM, from feature matching [4], view graph construction [42, 52], motion averaging [15] to bundle adjustment [44]. Yet enforcing the consistency is non-trivial in the learning-based settings. Instead of tweaking the network design or learning strategy, we seek a unified framework that effectively encodes geometries in different forms, and emphasize the efficacy of geometric reasoning for the remarkable improvement. Our contributions are summarized as follows:

- (1) We introduce traditional geometric quantities based on robust local descriptors into the learning pipeline, to complement the noisy photometric loss.
- (2) We propose a simple method to enforce pairwise and trinocular depth consistency in the unsupervised setting when both depth and pose are unknown.
- (3) Combined with a differentiable pixel selector mask, the proposed method outperform previous methods for the joint learning of depth and motion using monocular sequences.

## 2. Related Works

**Structure-from-Motion and Visual SLAM.** Structure-from-Motion (SfM) [2] and visual SLAM problems aim to simultaneously recover the camera pose and 3D structures from images. Both problems are well studied and render practical systems [9, 34, 50] by different communities for decades, with the latter emphasizes more on the real-time performance. The self-supervised depth and motion learning framework derives from direct SLAM methods [9, 10, 35]. Different from indirect methods [7, 25, 34] that use reliable sparse intermediate geometric quantities like local features [38], direct methods optimize the geometry using dense pixels in the image. With accurate photometric calibration such as gamma and vignetting correction [22], this formulation does not rely on sparse geometric computation and is able to generate finer-grained geometry. However, this formulation is less robust than indirect ones when the photometric loss is not meaningful, the scene containing moving or non-Lambertian objects.

**Supervised Approaches for Learning Depth.** Some early monocular depth estimation works rely on information from depth sensors [8, 39] without the aid of geometric relations. Liu *et al.* [28] combine deep CNN and conditional random field for estimating single monocular images. DeMoN [46] is an iterative supervised approach to jointly estimate optical flow, depth and motion. This coarse-to-fine process considers the use of stereopsis and produces good results with both depth and motion supervision.

**Unsupervised Depth Estimation from Stereo Match-**

**ing.** Based on warping-based view synthesis [55], Garg *et al.* [12] propose to learn depth using calibrated stereo camera pairs, in which per-pixel disparity is obtained by minimizing the image reconstruction loss. Godard *et al.* [14] improve this training paradigm with left-right consistency checking. Pilzer *et al.* [36] propose knowledge distillation from cycle-inconsistency refinement. These methods use synchronized and calibrated stereo images which are less affected by occlusion and photometric inconsistency. Therefore, this task is easier than ours which uses temporal multiview images and outputs relative poses in addition.

**Unsupervised Depth and Pose Estimation.** The joint unsupervised optimization of depth and pose starts from Zhou *et al.* [54] and Vijayanarasimhan *et al.* [47]. They propose similar approaches that use two CNNs to estimate depth and pose separately, and constrain the outcome with photometric loss. Later, a series of improvements [24, 31, 48, 51, 53] are proposed. Wang *et al.* [48] discuss the scale ambiguity and combine the estimated depth with direct methods [43, 9]. Zhan *et al.* [53] consider warping deep features from the neural nets instead of the raw pixel values. Klodt *et al.* [24] propose to integrate weak supervision from SfM methods. Mahjourian *et al.* [31] employ geometric constraints of the scene by enforcing an approximate ICP based matching loss. In this work, we follow the previous good practices, with the major distinction that we incorporate golden standards from indirect methods and enforce consistency terms to the state-of-the-art results.

## 3. Method

### 3.1. Problem Formulation

We first formalize the problem and present effective practices employed by previous methods [24, 31, 47, 48, 51, 53, 54]. Given adjacent  $N$ -view monocular image sequences (e.g.  $\{\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3\}$  for  $N = 3$ ), the unsupervised depth and motion estimation problem aims to simultaneously estimate the depth map  $\mathbf{D}_t$  of the target (center) image ( $\mathcal{I}_2$  in the 3-view case) and the 6-DoF relative poses  $\mathbf{T}_{t \rightarrow s} = [\mathbf{R}_{t \rightarrow s} | \mathbf{t}_{t \rightarrow s}] \in \mathcal{SE}(3)$  to  $N - 1$  source views ( $\mathcal{I}_1$  and  $\mathcal{I}_3$ ), using CNNs with photometric supervision.

For a source-target view pair  $(\mathcal{I}_s, \mathcal{I}_t)$ ,  $\mathcal{I}_t$  can be inversely warped to the source frame  $\mathcal{I}_s$  given the estimated depth map  $\mathbf{D}_t$  and the transformation from target to source  $\mathbf{T}_{t \rightarrow s}$ . Formally, given a pixel coordinate  $p_t$  in  $\mathcal{I}_t$  which is co-visible in  $\mathcal{I}_s$ , the pixel coordinate  $p_s$  in  $\mathcal{I}_s$  is given by the following equation which determines the warping transformation

$$p_s \sim \mathbf{K}_s [\mathbf{R}_{t \rightarrow s} | \mathbf{t}_{t \rightarrow s}] \mathbf{D}_t(p_t) \mathbf{K}_t^{-1} p_t, \quad (1)$$

where  $\sim$  denotes ‘equality in the homogeneous coordinates’,  $\mathbf{K}_s$  and  $\mathbf{K}_t$  are the intrinsics for the input image pair, and  $\mathbf{D}_t(p_t)$  is the depth for this particular  $p_t$  in  $\mathcal{I}_t$ .

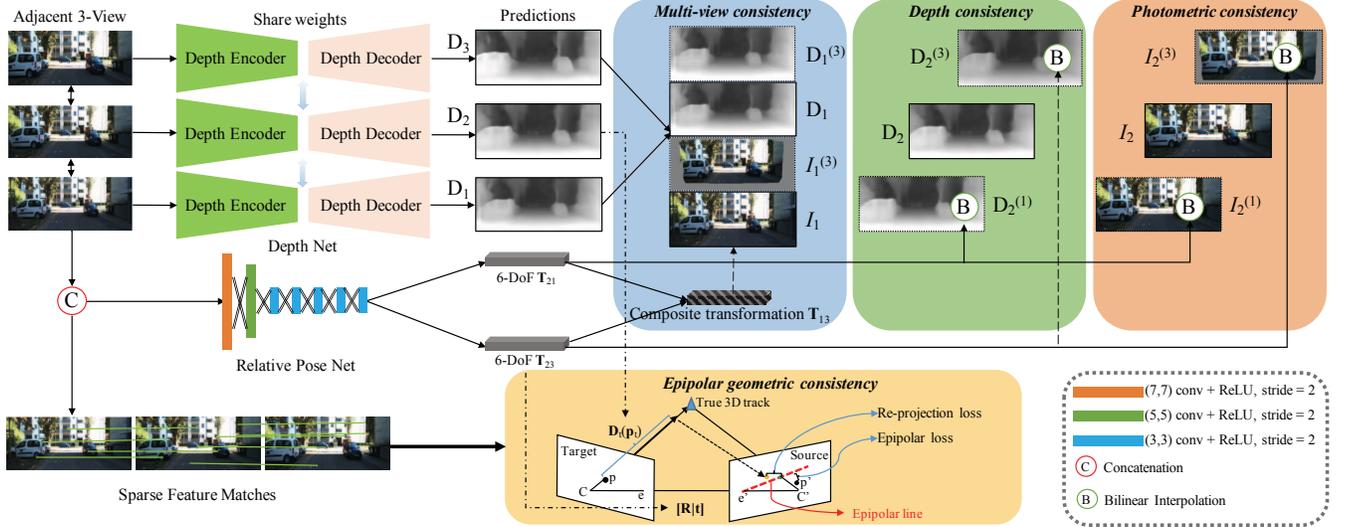


Figure 2. The architecture of our method. Besides photometric consistency (Section 3.1), we explore epipolar geometric consistency (Section 3.2), depth consistency and multi-view consistency (Section 3.3) to improve the depth and pose estimation.

With this coordinate transformation, synthesized images can be generated from the source view using the differentiable bilinear-sampling method [19]. The unsupervised framework then minimizes the pixel error between the target view and the synthesized image

$$\mathcal{L}_{pixel} = \frac{1}{|\mathcal{M}|} \sum_{\forall p_t \in \mathcal{M}} \left| \tilde{\mathcal{I}}_t^{(s)}(p_t | \mathbf{D}_t, \mathbf{T}_{t \rightarrow s}) - \mathcal{I}_t(p_t) \right|, \quad (2)$$

where  $\tilde{\mathcal{I}}_t^{(s)}$  represents the synthesized target image from source image.  $\mathcal{I}(p)$  is the function that maps the image coordinate  $p$  in image  $\mathcal{I}$  to pixel value, and the first term  $\tilde{\mathcal{I}}_t^{(s)}$  is the bilinear-sampling operation used to acquire the synthesized view given relative motion and depth.  $\mathcal{M}$  is a binary mask that determines if the inverse warping falls into a valid region in the source image, and can be computed analytically given the per-pixel depth and relative transformation.  $|\mathcal{M}|$  denotes the total number of valid pixels.

In addition to the per-pixel error, structured similarity (SSIM) [49] is shown to improve the performance [14, 51], which is defined on local image patches  $x$  and  $y$  rather than every single pixel. We follow the previous approaches [31, 51] to compute the SSIM loss on  $3 \times 3$  image patches ( $c_1 = 0.01^2$ ,  $c_2 = 0.03^2$ ) as follows

$$\mathcal{L}_{SSIM} = \frac{1}{2} \left[ 1 - \sum_{\forall x \in \tilde{\mathcal{I}}_t^{(s)}, \forall y \in \mathcal{I}_t} \frac{(2\mu_x \mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 \mu_y^2 + c_1)(\sigma_x + \sigma_y + c_2)} \right]. \quad (3)$$

The depth map is further constrained by the smoothness loss to push the gradients to propagate to nearby regions, known as the gradient locality issue [5]. Specifically, we adopt the image-aware smoothness formulation [14, 51]

which allows sharper depth changes on edge regions

$$\mathcal{L}_{smooth} = \sum_{\forall p_t \in \mathcal{I}_t} |\nabla \mathbf{D}_t(p_t)|^T \cdot e^{-|\nabla \mathcal{I}_t(p_t)|}, \quad (4)$$

where  $\nabla$  denotes the 2D differential operator for computing image gradients. Optimizing a combination of above loss terms wraps up the basic formulation of training objectives, which forms the baseline written as

$$\mathcal{L}_{baseline} = \alpha \mathcal{L}_{pixel} + (1 - \alpha) \mathcal{L}_{SSIM} + \beta \mathcal{L}_{smooth}. \quad (5)$$

However, there are drawbacks with the basic formulation. We then describe the key ingredients of our contributions.

### 3.2. Learning from Indirect Methods

The above view synthesis formulation requires several important assumptions: 1) the modeling scene should be static without moving objects; 2) the surfaces in the scene should be Lambertian; 3) no occlusion exists between adjacent views; 4) cameras should be photometrically calibrated, a technique adopted in direct SLAM methods [9, 10] to compensate for vignetting [22] and exposure time. Violation to any of the above criteria would lead to *photometric inconsistency*. The first three assumptions are inevitably violated to some extent because it is hard to capture temporally static images with no occlusion in the real world. The fourth restriction is often neglected by datasets with no photometric calibration parameters provided.

To address these limitations, previous methods [24, 54] additionally train a mask indicating whether the photometric loss is meaningful. Yet, we present a novel approach to tackle this issue by injecting indirect geometric information into the direct learning framework. Different from direct

methods that rely on dense photometric consistency, indirect methods for SfM and visual SLAM are based on sparse local descriptors such as SIFT [50] and ORB [34]. Local invariant features are much less likely to be affected by the scale and illumination changes and can be implicitly embedded into the learning framework.

**Symmetric epipolar error.** Assuming the pinhole camera model, the feature matches  $\mathcal{S}_{t \leftrightarrow s} = \{\mathbf{p} \leftrightarrow \mathbf{p}'\}$  between the target and source views satisfy the epipolar constraint, where  $\mathbf{p}$  and  $\mathbf{p}'$  are the calibrated image coordinates. The loss with the feature matches and the estimated pose can be quantified using the *symmetric epipolar distance* [16]

$$\mathcal{L}_{epi}(\mathcal{S}|\mathbf{R}, \mathbf{t}) = \sum_{\forall (\mathbf{p}, \mathbf{p}') \in \mathcal{S}} \left( \frac{\mathbf{p}^T \mathbf{E} \mathbf{p}}{\sqrt{(\mathbf{E} \mathbf{p})_{(1)}^2 + (\mathbf{E} \mathbf{p})_{(2)}^2}} + \frac{\mathbf{p}'^T \mathbf{E} \mathbf{p}'}{\sqrt{(\mathbf{E} \mathbf{p}')_{(1)}^2 + (\mathbf{E} \mathbf{p}')_{(2)}^2}} \right), \quad (6)$$

where  $\mathbf{E}$  being the essential matrix computed by  $\mathbf{E} = [\mathbf{t}]_{\times} \mathbf{R}$ ,  $[\cdot]_{\times}$  is the matrix representation of the cross product with  $\mathbf{t}$ . We simply omit the subindices for conciseness ( $\mathcal{S}$  for  $\mathcal{S}_{t \leftrightarrow s}$ ,  $\mathbf{R}$  for  $\mathbf{R}_{t \rightarrow s}$ ,  $\mathbf{t}$  for  $\mathbf{t}_{t \rightarrow s}$ ).

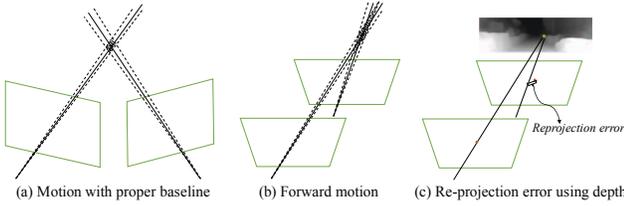


Figure 3. (a) For two images with proper motion baseline, the uncertainty (shaded region) is small. (b) For forward motion with narrow baseline, the uncertainty is large. (c) The re-projection error unites estimated depth and pose with sparse features, and does not involve triangulation uncertainty.

**Re-projection error.** The epipolar constraint does not concern the depth in its formulation. To involve depth optimization using the feature match supervision, there are generally two methods: 1) triangulate the correspondence  $\mathbf{p} \leftrightarrow \mathbf{p}'$  using the optimal triangulation method [16] assuming the Gaussian noise model, to obtain the 3D track for depth supervision; 2) back-project 2D features in one image using the estimated depth to compute the 3D track, and re-project the 3D track to another image to compute the re-projection error. We take the second method because the estimated depth and pose are sufficient to compute the 3D loss, and triangulation is often imprecise for ego-motion driving scenes [6, 13] (see Figure 3 for illustration, and a mathematically rigorous explanation in the Appendix).

$$\mathcal{L}_{reproj}(\mathcal{S}|\mathbf{R}, \mathbf{t}, \mathbf{D}_t) = \sum_{\forall \mathbf{p} \leftrightarrow \mathbf{p}' \in \mathcal{S}} \left\| [\mathbf{R}|\mathbf{t}] \hat{\mathbf{D}}_t(\mathbf{p}) \mathbf{p} - \mathbf{p}' \right\|_2, \quad (7)$$

where  $\hat{\mathbf{D}}_t(\mathbf{p})$  is the bilinear-sampling operation [19] in the target depth map as the feature coordinate  $p$  is not an integer. Minimizing re-projection error using feature matches can be viewed as creating sparse anchors between the weak geometric supervision and the estimated depth and pose. In contrast, Equation 6 does not involve the estimated depth.

Since outliers may exist if they lie close to the epipolar line, we use the pairwise matches that are confirmed in three views [17]. Minimizing the epipolar and re-projection errors of all matches using CNNs mimics the non-linear pose estimation [3]. The experiment shows that this weak supervisory signal significantly improves the pose estimation and is superior to other SfM supervisions such as [24].

### 3.3. Consistent Depth Estimation

In this section, we describe the depth estimation module. Previous methods, whether operating on three or five views, are pairwise approaches in essence because loss terms are computed pairwise from the source frame to the target frame. Even though the pose network outputs  $N - 1$  relative poses at once, it is unknown if these relative poses are aligned to the same scale. We propose the motion-consistent depth estimation formulation to address this issue. Rather than minimizing the loss between the target frame and adjacent source frames, our proposed formulation also considers the depth and motion consistency between adjacent frames themselves.

**Forward-backward consistency.** As shown in Figure 2, our network architecture estimates the depth maps of the target image ( $\mathcal{I}_t$ ), as well as the forward and backward depths. Inspired by [14, 37] that uses left-right consistency on stereo images, we propose *forward-backward consistency* for monocular images. In addition to bilinear-sampling pixel values, it samples the estimated depth maps of forward and backward images ( $\mathcal{I}_s$ ). This process generates two synthesized depth maps  $\hat{\mathbf{D}}_t^{(s)}$  that can be used to constrain the estimation of the target image depth map  $\mathbf{D}_t$ .

However, the availability of only monocular images makes the problem more challenging. For learning with stereo images, the images are rectified in advance so the scale ambiguity issue is not considered. While for learning monocular depth, the estimated depth is determined only up to scale, therefore the alignment of depth scale is necessary before constraining the depth discrepancy. We first normalize the target depth map using its mean  $\mathbf{D}_t := \mathbf{D}_t / \text{mean}(\mathbf{D}_t)$  to resolve the scale ambiguity of the target depth [48], which determines the scale of relative poses. Then we apply a mean alignment to the synthesized depth maps and the normalized target depth map in the corresponding region informed by the analytical mask  $\mathcal{M}$  (Equa-



Figure 4. Differentiable mask composition. The inconsistency incurred by a moving object is filtered while the meaningful loss is preserved.

tion 2), and further optimize the depth discrepancy

$$\mathcal{L}_{depth} = \frac{1}{|\mathcal{M}|} \sum_{\forall p \in \mathcal{M}} \left| \frac{mean(\mathbf{D}_t \circ \mathcal{M})}{mean(\tilde{\mathbf{D}}_t^{(s)} \circ \mathcal{M})} \cdot \tilde{\mathbf{D}}_t^{(s)}(p) - \mathbf{D}_t(p) \right|, \quad (8)$$

where  $\circ$  means the element-wise multiplication and the loss is averaged over all the valid pixel  $p$  in the mask  $\mathcal{M}$ .

**Multi-view consistency.** The above losses are all defined on the single target image (e.g. smoothness loss) or among image pairs, even though the input is  $N$ -view ( $N \geq 3$ ) image sequences. The pose network outputs  $N - 1$  relative poses between the target and source images, but the  $N - 1$  relative poses are only weakly connected by the monocular depth. To strengthen the scale consistency for triplet relation, we propose the multi-view consistency loss which penalizes inconsistency of the forward depth and backward depth using the target image as a bridge for scale alignment. Formally, given image sequence  $(\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3)$  with  $\mathcal{I}_2$  the target image, and corresponding pose and depth predictions  $(\mathbf{T}_{2 \rightarrow 1}, \mathbf{T}_{2 \rightarrow 3})$  and  $(\mathbf{D}_1, \mathbf{D}_2, \mathbf{D}_3)$ , we again obtained the normalized depth map  $\bar{\mathbf{D}}_1 = s_{12} \cdot \mathbf{D}_1$  where the scaling ratio  $s_{12} = \frac{mean(\mathbf{D}_2 \circ \mathcal{M}_{12})}{mean(\bar{\mathbf{D}}_2^{(1)} \circ \mathcal{M}_{12})}$  as used in Equation 8. The transformation from the backward image  $\mathcal{I}_1$  to the forward image  $\mathcal{I}_3$  is  $\mathbf{T}_{1 \rightarrow 3} = \mathbf{T}_{2 \rightarrow 1}^{-1} \cdot \mathbf{T}_{2 \rightarrow 3}$ . The multiview loss minimizes the depth consistency term and photometric consistency term as

$$\begin{aligned} \mathcal{L}_{multi} = & \alpha \mathcal{L}_{pixel}(\mathcal{I}_1, \tilde{\mathcal{I}}_1^{(3)}) + (1 - \alpha) \mathcal{L}_{SSIM}(\mathcal{I}_1, \tilde{\mathcal{I}}_1^{(3)}) \\ & + \frac{1}{|\mathcal{M}_{13}|} \sum_{\forall p \in \mathcal{M}_{13}} \left| \bar{\mathbf{D}}_1(p) - \bar{\mathbf{D}}_1^{(3)}(p) \right|, \quad (9) \end{aligned}$$

where  $\tilde{\mathcal{I}}_1^{(3)}$  and  $\bar{\mathbf{D}}_1^{(3)}$  are the synthesized image and synthesized normalized depth given  $\bar{\mathbf{D}}_3$  and  $\mathbf{T}_{1 \rightarrow 3}$ . The sub-indices 1 and 3 are interchangeable in the above Equation 9.  $\mathcal{L}_{multi}$  goes beyond the pairwise loss terms  $\mathcal{L}_{pixel}$ ,  $\mathcal{L}_{SSIM}$ ,  $\mathcal{L}_{epi}$  and  $\mathcal{L}_{depth}$  because it utilizes the chained pose and pushes the two relative poses to be aligned on the same scale. This benefits monocular SLAM because it facilitates the incremental localization by aligning multiple  $N$ -view outputs, as we show in Section 4.4.

### 3.4. Differentiable Sparse Feature Selection

Photometric inconsistency inevitably exists due to occlusion or non-Lambertian properties. Previous works employ

an additional branch to regress an uncertainty map, which helps a little [54]. Instead, we follow the explicit occlusion modeling approach [41] which does not rely on the data-driven uncertainty. We have observed that photometric inconsistency such as moving objects usually incurs larger photometric errors (Figure 4(b)). On the other hand, image region with small gradient changes does not offer meaningful supervisory information because of the gradient locality issue [5] (Figure 4(c)).

Therefore, we combine the error mask with the gradient mask to select the meaningful sparse features, inspired by the direct sparse odometry [9] but can be fit into the differentiable training pipeline. Given the pixel error map, we compute the error histogram and mask out the pixels which are above  $\sigma_e (= 90)$ -th percentile. We also compute the gradient mask and keep only the values that are above  $\sigma_g (= 90)$ -th percentile. The final composite mask is the multiplication of both masks with dynamic thresholding. As shown in Figure 4(d), this mask operation filters out a majority of photometric inconsistency regions like the moving car. The composite mask is only used for the final depth refinement when the error suppression mask is stable, otherwise we observe a performance drop if training from scratch.

Our final formulation takes into account the basic losses in Equation 5, the geometric terms, as well as the consistency terms, written as

$$\begin{aligned} \mathcal{L}_{total} = & \alpha \mathcal{L}_{pixel} + (1 - \alpha) \mathcal{L}_{SSIM} + \beta \mathcal{L}_{smooth} + \\ & \gamma_1 \mathcal{L}_{epi} + \gamma_2 \mathcal{L}_{reproj} + \mu_1 \mathcal{L}_{depth} + \mu_2 \mathcal{L}_{multi}. \quad (10) \end{aligned}$$

The weighting for different losses are set empirically given hyper-parameters in previous methods and our attempts ( $\alpha = 0.15, \beta = 0.1, \gamma_1 = \gamma_2 = 0.001, \mu_1 = \mu_2 = 0.1$ ). We also try to learn the optimal weighting using homoscedastic uncertainty [21], but find no better result than empirically setting the weights.

## 4. Experiments

### 4.1. Training Dataset

**KITTI.** We evaluate our method on the KITTI datasets [13, 32], using the raw dataset with Eigen split [8] for depth estimation, and the odometry dataset for pose estimation. Images are down-sampled to  $128 \times 416$  to facilitate the training and provide a fair evaluation setting. For Eigen split,

Method	Supervision	Dataset	Cap (m)	Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Eigen <i>et al.</i> [8] Fine	Depth	K	80	0.203	1.548	6.307	0.282	0.702	0.890	0.958
Liu <i>et al.</i> [28]	Depth	K	80	0.202	1.614	6.523	0.275	0.678	0.895	0.965
Godard <i>et al.</i> [14]*	Stereo/Pose	K	80	0.148	1.344	5.927	0.247	0.803	0.922	0.964
Godard <i>et al.</i> [14]*	Stereo/Pose	K + CS	80	<b>0.114</b>	<b>0.898</b>	<b>4.935</b>	<b>0.206</b>	<b>0.861</b>	<b>0.949</b>	<b>0.976</b>
Zhou <i>et al.</i> [54] updated	No	K	80	0.183	1.595	6.709	0.270	0.734	0.902	0.959
Zhou <i>et al.</i> [54] updated	No	K	-	0.185	2.170	6.999	0.271	0.734	0.901	0.959
Klodt <i>et al.</i> [24]	No	K	80	0.166	1.490	5.998	-	0.778	0.919	0.966
Mahjourian <i>et al.</i> [31]	No	K	80	0.163	1.24	6.22	0.25	0.762	0.916	0.968
Wang <i>et al.</i> [48]	No	K	80	0.151	1.257	5.583	0.228	0.810	0.936	0.974
Yin <i>et al.</i> [51]	No	K	80	0.155	1.296	5.857	0.233	0.793	0.931	0.973
Yin <i>et al.</i> [51]	No	K	-	0.156	1.470	6.197	0.235	0.793	0.931	0.972
Yin <i>et al.</i> [51] updated	No	K + CS	80	0.149	1.060	5.567	0.226	0.796	0.935	0.975
Ours	No	K	80	0.140	1.025	5.394	0.222	0.816	0.938	0.974
Ours	No	K	-	0.140	1.026	5.397	0.222	0.816	0.937	0.974
Ours	No	K + CS	80	<b>0.139</b>	<b>0.964</b>	<b>5.309</b>	<b>0.215</b>	<b>0.818</b>	<b>0.941</b>	<b>0.977</b>
Garg <i>et al.</i> [12]	Stereo/Pose	K	50	0.169	1.080	5.104	0.273	0.740	0.904	0.962
Zhou <i>et al.</i> [54]	No	K	50	0.201	1.391	5.181	0.264	0.696	0.900	0.966
Yin <i>et al.</i> [51]	No	K + CS	50	0.147	0.936	4.348	0.218	0.810	0.941	0.977
Ours	No	K	50	<b>0.133</b>	<b>0.778</b>	<b>4.069</b>	<b>0.207</b>	<b>0.834</b>	<b>0.947</b>	<b>0.978</b>

Table 1. Single-view depth estimation performance. The statistics for the compared methods are excerpted from corresponding papers, except that the results marked with ‘updated’ are captured from the websites. ‘K’ represents KITTI raw dataset (Eigen split) and CS represents cityscapes training dataset. The method [14] marked with \* are trained and tested on larger scale ( $256 \times 512$ ) images, whereas others use  $128 \times 416$  images. ‘-’ in Cap(m) means no maximum depth filtering is applied. The metrics marked by red means ‘the lower the better’ and the ones marked by green means ‘the higher the better’. The best results for each category are bolded.

we use 20129 images for training and 2214 images for validation. The 697 testing images are selected by [8] from 28 scenes whose images are excluded from the training set. For the KITTI odometry dataset, we follow the previous convention [51, 54] to train the model on sequence 00-08 and test on sequence 09-10. We further split sequence 00-08 to 18361 images for training and 2030 for validation.

**Cityscapes.** We also try pre-training the model on the Cityscapes [6] dataset since starting from a pre-trained model boosts the performance [54]. The process is conducted without adding feature matches for 60k steps. 88084 images are used for training and 9659 images for validation.

## 4.2. Implementation Details

**Data preparation.** We extract SIFT [30] feature matches as the weak geometric supervision using SiftGPU [50] offline. The putative matches are further filtered by geometric verification [17] with RANSAC [11]. 100 feature pairs are randomly sampled and used for training. Matches are only used for training and not necessary for inference.

**Learning.** We implement our pipeline using Tensorflow [1]. The depth estimation part follows [51] which uses ResNet-50 [18] as the depth encoder. The relative pose net follows [54, 51] which is a 7-layer CNN, with the lengths of feature maps reduced by half and the number of feature channels multiplied by two from each previous layer. If not explicitly specified, we train the neural networks using 3-view image sequences as the photometric error would accumulate for longer input sequences. We use the Adam [23]

solver with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , a learning rate of 0.0001 and a batch size of 4.

**Training efficiency.** The proposed method takes longer time per step due to more depth estimations and loss computations. With a single GTX 1080 Ti, training takes 0.35s per step compared with 0.19s for the baseline approach based on Equation 5. It is noted that the inference efficiency is the same as the baseline.

## 4.3. Depth Estimation

The evaluation of depth estimation follows previous works [31, 51, 54]. As shown in Table 1, our method achieves the best performance among all unsupervised methods that jointly learn depth and pose. Previous methods often filter the predicted depth map by setting a maximum depth at 50m or 80m (the ground-truth depth range is within 80m) before computing depth error, since distant pixels may have prediction outliers. We also evaluate the performance without this filtering step, marked by ‘-’ in the Cap(m) column. It shows that without capping the maximum depth, [51, 54] become worse while our result seldom changes, meaning our consistent training renders depth predictions with little noise. Figure 5 provides a qualitative comparison of the predictions. We show the depth value (the nearer the darker) instead of the inverse depth (disparity) parameterization, which highlights the distant areas.

Since both Klodt *et al.* [24] and ours use self-supervised weak supervisions, we redo the experiments in [24] that use self-generated poses and sparse depth maps from ORB-

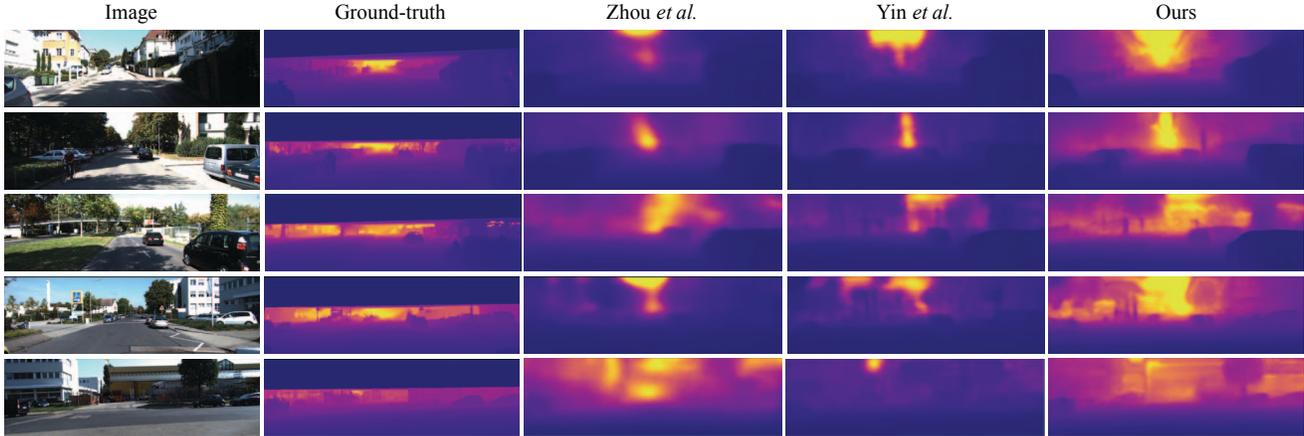


Figure 5. Qualitative comparison for depth estimation on the Eigen split. The predicted depth maps are first aligned with the ground-truth using mean. Then the depth values larger than 80m are set to 80m to ensure a consistent color scale. It shows that our result best reflects the ground-truth depth range and contains richer details (best view in color).

Method	Seq 09	Seq 10
ORB-SLAM2 [34]	$0.014 \pm 0.008$	$0.012 \pm 0.011$
Zhou <i>et al.</i> [54] updated (5-frame)	$0.016 \pm 0.009$	$0.013 \pm 0.009$
Yin <i>et al.</i> [51] (5-frame)	$0.012 \pm 0.007$	$0.012 \pm 0.009$
Mahjourian <i>et al.</i> [31], no ICP (3-frame)	$0.014 \pm 0.010$	$0.013 \pm 0.011$
Mahjourian <i>et al.</i> [31], with ICP (3-frame)	$0.013 \pm 0.010$	$0.012 \pm 0.011$
Klodt <i>et al.</i> [24] (5-frame)	$0.014 \pm 0.007$	$0.013 \pm 0.009$
Ours <i>et al.</i> (3-frame)	<b><math>0.009 \pm 0.005</math></b>	<b><math>0.008 \pm 0.007</math></b>

Table 2. Pose estimation evaluation. All the learning-based methods are trained and tested on  $128 \times 416$  images, while ORB-SLAM2 are tested on full-sized ( $370 \times 1226$ ) images.

SLAM2 [34] for weak supervision, fixing other settings. We obtain slightly better statistics which still lags behind the proposed method that uses feature matches. This implies that the raw matches are more robust as the supervisory signal, whereas using pose and depth computed from SfM/SLAM is possible to introduce additional bias inherited from the PnP [27] or triangulation algorithms.

#### 4.4. Pose Estimation

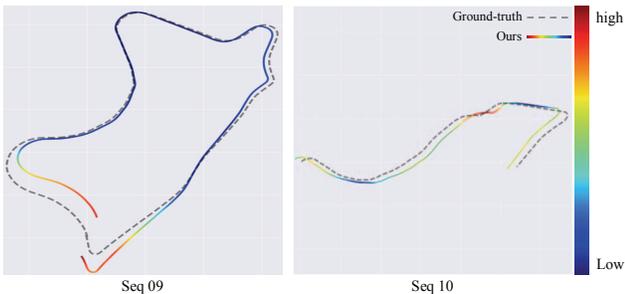


Figure 6. The chained trajectory of the proposed method drawn with the ground-truth on KITTI sequence 09/10. The color bar on the right shows the scale of alignment error.

We evaluate the performance of relative pose estimation on the KITTI odometry dataset. We have observed that with the pairwise matching supervision, the result for motion estimation has been extensively improved. We measure the Absolute Trajectory Error (ATE) over  $N$ -frame snippets. The mean error and variance are averaged from the full sequence. As shown in Table 2, with the same underlying network structure, the proposed method outperforms state-of-the-art methods by a large margin.

However, the above comparison is in favor of learning-based approaches which are only able to generate  $N$ -view pose segments, but not fair for mature SLAM systems which emphasizes the accuracy of the full trajectory. To demonstrate that our method produces consistent pose estimation, we chain the relative poses by averaging the two overlapping frames of 3-view snippets. We first align the chained motion with the ground-truth trajectory by estimating a similarity transformation [45], and then compute the average of APE for each frame. As shown in Figure 6, even without global motion averaging techniques [15], our method achieves comparable performance (8.82m/23.09m median APE for Seq. 9/10) against monocular ORB-SLAM [33] (36.83m/5.74m median APE for Seq. 9/Seq. 10) without loop closure. This comparison just provides a fair comparison in terms of the full sequence, yet by no means shows the learning-based method has surpassed tradition VO methods. In fact, monocular ORB-SLAM with loop closure and global bundle adjustment results in a much smaller 7.08m median APE for Seq. 9 (Seq. 10 stays unchanged because it has no loop).

#### 4.5. Ablation Study

**Performance with different modules.** We conduct an ablation study to show the effect of each component. The models for depth and pose evaluation are trained solely on

Loss Configuration						Depth (KITTI raw Eigen split)							Pose (KITTI odometry)	
Baseline	Epipolar	Re-projection	Forward-backward	Multi-view	Mask	Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25^3$	$\delta < 1.25^2$	$\delta < 1.25^3$	Seq 09	Seq 10
✓	-	-	-	-	-	0.163	1.371	6.275	0.249	0.773	0.918	0.966	0.014 ± 0.009	0.012 ± 0.012
✓	✓	-	-	-	-	0.159	1.287	5.725	0.239	0.791	0.927	0.969	0.010 ± 0.005	0.009 ± 0.008
✓	✓	-	-	-	-	0.152	1.205	5.56	0.227	0.800	0.935	0.973	0.009 ± 0.005	0.009 ± 0.008
✓	✓	✓	✓	-	-	0.146	1.391	5.791	0.229	0.814	0.936	0.972	0.009 ± 0.005	0.008 ± 0.007
✓	✓	✓	✓	✓	-	0.143	1.114	5.681	0.225	0.816	0.938	0.974	<b>0.009 ± 0.005</b>	<b>0.008 ± 0.007</b>
✓	✓	✓	✓	✓	✓	<b>0.140</b>	<b>1.025</b>	<b>5.394</b>	<b>0.222</b>	<b>0.816</b>	<b>0.938</b>	<b>0.974</b>	<b>0.009 ± 0.005</b>	<b>0.008 ± 0.007</b>
✓ (5-view)	-	-	-	-	-	0.169	1.607	6.129	0.255	0.779	0.917	0.963	0.014 ± 0.009	0.013 ± 0.009
✓ (5-view)	✓	✓	-	-	-	0.157	1.449	5.796	0.239	0.803	0.929	0.970	0.012 ± 0.008	0.010 ± 0.007

Table 3. Evaluation of different training loss configurations. All models are either solely trained on KITTI raw dataset (for depth) or KITTI odometry dataset (for pose) without pre-training on Cityscapes. The depth estimation performance is evaluated with maximum depth set/capped at 80m. All models except the last two are trained on 3-view image sequences. The best result for each metric is **bolded**.

KITTI raw dataset and odometry dataset respectively. We choose an incremental order for the proposed techniques to avoid too many loss term combinations. As shown in Table 3, we have the following observations:

- The re-implemented baseline model, using Equation 5, has already surpassed several models [24, 31, 54]. The reasons can be attributed to the more capable depth encoder ResNet-50, which is also used by [51].
- The result for pose estimation is greatly improved with the epipolar loss term  $\mathcal{L}_{epi}$ . It shows the efficacy of using raw feature matches as the weakly supervised signal. However, the improvement for depth estimation is not as significant as pose estimation.
- Re-projection loss further improves the depth inference. The improvements for pose estimation brought by ingredients other than the epipolar loss are marginal.
- The depth consistency and multi-view consistency are the essential parts for the improvement in depth estimation.

In summary, the epipolar geometric supervision helps the pose estimation most, while the geometric consistency terms in Section 3.3 essentially improve depth estimation.

**Sequence Length.** The multi-view depth consistency loss boosts the depth estimation. However, the performance boost can be also attributed to using longer image snippets, since similar second-order relations can be exploited by using 5-view image sequences for training. Therefore, we further evaluate the performance of using 5-view images. As shown in Table 3, training on longer image sequences would deteriorate the performance, because long sequences also contain larger photometric noises. It shows that the proposed formulation elevates the results not from more data, but the consistency embedded in geometric relations.

#### 4.6. Generalization on Make3D

To illustrate that the proposed method is able to generalize to other datasets unseen in the training, we compare to several supervised/self-supervised methods on the Make3D dataset [40], using the same evaluation protocol as in [14]. As shown in Table 4, our best model achieves reasonable generalization ability and even beats several super-

Method	Supervision		Metrics			
	depth	pose	Abs Rel	Sq Rel	RMSE	RMSE log <sub>10</sub>
Karsch <i>et al.</i> [20]†	✓	-	0.417	4.894	8.172	0.144
Liu <i>et al.</i> [29]†	✓	-	0.462	6.625	9.972	0.161
Laina <i>et al.</i> [26] †	✓	-	0.198	1.665	5.461	0.082
Godard <i>et al.</i> [14]	-	✓	0.443	7.112	8.860	0.142
Zhou <i>et al.</i> [54]	-	-	0.392	4.473	8.307	0.194
Ours	-	-	0.378	4.348	7.901	0.183

Table 4. Generalization experiments on Make3D. The evaluation metrics are the same as the ones in Table 1 except the last one (RMSE log<sub>10</sub>) to conform with [20]. The methods marked with † are trained on Make3D. The depth estimation is evaluated with maximum depth capped at 70m. We use the center-cropped images as in [14] and resize them to 128 × 416 for inference.

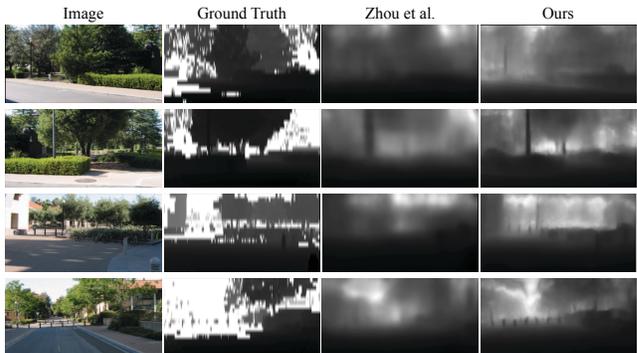


Figure 7. Sample depth predictions on the Make3D dataset. Both our method and SfMlearner [54] are trained on Cityscapes+KITTI.

vised methods on some metrics. A qualitative comparison is shown in Figure 7.

## 5. Conclusion

We have presented an unsupervised pose and depth estimation pipeline that absorbs both the geometric principles and learning-based metrics. We emphasize on the consistency issue and propose novel ingredients to make the result more robust and reliable. Yet, we should realize that the current learning-based methods are still far from solving the SfM problem in an end-to-end fashion. Further investigations include enforcing consistency across the whole dataset, such as incorporating loop closure and bundle adjustment techniques into the learning-based methods.

## References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: a system for large-scale machine learning. In *OSDI*, 2016. 6
- [2] Sameer Agarwal, Noah Snavely, Ian Simon, Steven M Seitz, and Richard Szeliski. Building rome in a day. In *ICCV*, 2009. 2
- [3] Adrien Bartoli and Peter Sturm. Non-linear estimation of the fundamental matrix with minimal parameters. *PAMI*, 2004. 4
- [4] Adam Baumberg. Reliable feature matching across widely separated views. In *CVPR*, 2000. 2
- [5] James R Bergen, Patrick Anandan, Keith J Hanna, and Ramesh Hingorani. Hierarchical model-based motion estimation. In *ECCV*, 1992. 1, 3, 5
- [6] 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*, 2016. 1, 4, 6
- [7] Andrew J Davison, Ian D Reid, Nicholas D Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *PAMI*, 2007. 2
- [8] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NIPS*, 2014. 2, 5, 6, 11
- [9] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *PAMI*, 2017. 2, 3, 5
- [10] Jakob Engel, Thomas Schöps, and Daniel Cremers. Lsdslam: Large-scale direct monocular slam. In *ECCV*, 2014. 2, 3
- [11] 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*, 1981. 6
- [12] Ravi Garg, Vijay Kumar BG, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *ECCV*, 2016. 2, 6, 11
- [13] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013. 1, 4, 5
- [14] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017. 2, 3, 4, 6, 8, 11
- [15] Venu Madhav Govindu. Robustness in motion averaging. In *ACCV*, 2006. 2, 7, 11
- [16] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 4, 11
- [17] Richard I Hartley. In defense of the eight-point algorithm. *PAMI*, 1997. 4, 6
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [19] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *NIPS*, 2015. 1, 3, 4
- [20] Kevin Karsch, Ce Liu, and Sing Bing Kang. Depth transfer: Depth extraction from video using non-parametric sampling. *PAMI*, 2014. 8
- [21] Alex Kendall, Roberto Cipolla, et al. Geometric loss functions for camera pose regression with deep learning. In *CVPR*, 2017. 5
- [22] Seon Joo Kim and Marc Pollefeys. Robust radiometric calibration and vignetting correction. *PAMI*, 2008. 2, 3
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [24] Maria Klodt and Andrea Vedaldi. Supervising the new with the old: learning sfm from sfm. In *ECCV*, 2018. 2, 3, 4, 6, 7, 8
- [25] Kurt Konolige and Motilal Agrawal. Frameslam: From bundle adjustment to real-time visual mapping. *IEEE Transactions on Robotics*, 2008. 2
- [26] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *3DV*, 2016. 8
- [27] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnnp: An accurate o(n) solution to the pnp problem. *IJCV*, 2009. 7
- [28] Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian D Reid. Learning depth from single monocular images using deep convolutional neural fields. *PAMI*, 2016. 2, 6, 11
- [29] Miaomiao Liu, Mathieu Salzmann, and Xuming He. Discrete-continuous depth estimation from a single image. In *CVPR*, 2014. 8
- [30] David G Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 2004. 6
- [31] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In *CVPR*, 2018. 2, 3, 6, 7, 8
- [32] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *CVPR*, 2015. 5
- [33] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE Transactions on Robotics*, 2015. 7, 11
- [34] Raul Mur-Artal and Juan D Tardós. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE Transactions on Robotics*, 2017. 2, 4, 7
- [35] Richard A Newcombe, Steven J Lovegrove, and Andrew J Davison. Dtam: Dense tracking and mapping in real-time. In *ICCV*, 2011. 2
- [36] Andrea Pilzer, Stéphane Lathuilière, Nicu Sebe, and Elisa Ricci. Refine and distill: Exploiting cycle-inconsistency and knowledge distillation for unsupervised monocular depth estimation. In *CVPR*, 2019. 2
- [37] Matteo Poggi, Fabio Tosi, and Stefano Mattoccia. Learning monocular depth estimation with unsupervised trinocular assumptions. In *3DV*, 2018. 4
- [38] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *ICCV*, 2011. 2

- [39] Ashutosh Saxena, Sung H Chung, and Andrew Y Ng. Learning depth from single monocular images. In *NIPS*, 2006. 2
- [40] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Make3d: Learning 3d scene structure from a single still image. *PAMI*, 2009. 8
- [41] Tianwei Shen, Zixin Luo, Lei Zhou, Hanyu Deng, Runze Zhang, Tian Fang, and Long Quan. Beyond photometric loss for self-supervised ego-motion estimation. In *ICRA*, 2019. 5
- [42] Tianwei Shen, Siyu Zhu, Tian Fang, Runze Zhang, and Long Quan. Graph-based consistent matching for structure-from-motion. In *ECCV*, 2016. 2
- [43] Frank Steinbrücker, Jürgen Sturm, and Daniel Cremers. Real-time visual odometry from dense rgb-d images. In *ICCV Workshops*, 2011. 2
- [44] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment: modern synthesis. In *International workshop on vision algorithms*, 1999. 2
- [45] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *PAMI*, 1991. 7
- [46] Benjamin Ummenhofer, Huizhong Zhou, Jonas Uhrig, Nikolaus Mayer, Eddy Ilg, Alexey Dosovitskiy, and Thomas Brox. Demon: Depth and motion network for learning monocular stereo. In *CVPR*, 2017. 2
- [47] Sudheendra Vijayanarasimhan, Susanna Ricco, Cordelia Schmid, Rahul Sukthankar, and Katerina Fragkiadaki. Sfmnet: Learning of structure and motion from video. *arXiv preprint arXiv:1704.07804*, 2017. 1, 2
- [48] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *CVPR*, 2018. 2, 4, 6
- [49] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 2004. 3
- [50] Changchang Wu et al. Visualsfm: A visual structure from motion system. 2011. 2, 4, 6
- [51] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *CVPR*, 2018. 1, 2, 3, 6, 7, 8
- [52] Christopher Zach, Manfred Klopschitz, and Marc Pollefeys. Disambiguating visual relations using loop constraints. In *CVPR*, 2010. 2
- [53] Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, and Ian Reid. Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In *CVPR*, 2018. 2
- [54] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017. 1, 2, 3, 5, 6, 7, 8
- [55] C Lawrence Zitnick, Sing Bing Kang, Matthew Uyttendaele, Simon Winder, and Richard Szeliski. High-quality video view interpolation using a layered representation. In *TOG*, 2004. 2