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

Mono-SF: Multi-View Geometry Meets Single-View Depth for Monocular Scene Flow Estimation of Dynamic Traffic Scenes

Fabian Brickwedde^{1,2} Steffen Abraham¹ Rudolf Mester^{3,2} ¹ Robert Bosch GmbH, Hildesheim, Germany ² VSI Lab, CS Dept., Goethe University, Frankfurt, Germany ³ Norwegian Open AI Lab, CS Dept. (IDI), NTNU Trondheim, Norway {Fabian.Brickwedde; Steffen.Abraham}@de.bosch.com

Abstract

Existing 3D scene flow estimation methods provide the 3D geometry and 3D motion of a scene and gain a lot of interest, for example in the context of autonomous driving. These methods are traditionally based on a temporal series of stereo images. In this paper, we propose a novel monocular 3D scene flow estimation method, called Mono-SF. Mono-SF jointly estimates the 3D structure and motion of the scene by combining multi-view geometry and singleview depth information. Mono-SF considers that the scene flow should be consistent in terms of warping the reference image in the consecutive image based on the principles of multi-view geometry. For integrating single-view depth in a statistical manner, a convolutional neural network, called ProbDepthNet, is proposed. ProbDepthNet estimates pixelwise depth distributions from a single image rather than single depth values. Additionally, as part of ProbDepth-Net, a novel recalibration technique for regression problems is proposed to ensure well-calibrated distributions. Our experiments show that Mono-SF outperforms state-of-theart monocular baselines and ablation studies support the Mono-SF approach and ProbDepthNet design.

1. Introduction

In applications such as mobile robots or autonomous vehicles a representation of the surrounding environment is utilized, e.g. to fulfill a navigation task. From a computer vision point of view, the 3D position and motion of a pixel in the image is denoted as *3D scene flow* [59,60], which is traditionally estimated based on a temporal series of stereo images [4, 44, 61]. In this work, we propose a novel scene flow estimation method, *Mono-SF*, for a monocular camera setup focusing on dynamic traffic scenes. Monocular camera systems are often preferred over stereo cameras due to being more cost efficient and to avoid the effort of calibrat-



Figure 1. Overview of Mono-SF for monocular scene flow estimation. Mono-SF jointly optimizes the 3D geometry of a set of planes with the 6D motion of rigid bodies considering a) a photometric distance by warping the reference image into the consecutive image, b) probabilistic depth distributions provided by ProbDepthNet and c) scene model smoothness priors.

ing the stereo rig. However, 3D scene flow estimation is an ill-posed problem in a monocular camera setup. To solve the ambiguity, previous monocular methods assumed that the moving objects are in contact with the surrounding environment [6, 8, 51] or that the scene follows a smoothness prior regarding surface and motion [36, 46, 66]. These assumptions might be violated and the methods still require a relative translational motion of the camera to the scene. In contrast to the multi-view geometry-based approaches, methods were proposed (e.g. [10, 14, 19]) that provide depth estimates from a single image at a reasonable level of quality. However, single-view depth estimation and multi-view geometry are mostly tackled as two individual tasks or fused in a way that is only applicable for static scenes [12,54,71]. Our proposed Mono-SF method combines multi-view geometry with single-view depth information in a probabilistic optimization framework to provide consistent 3D scene flow estimates. Thereby, both kinds of information are exploited and the single-view depth serves to solve the multiview geometry-based ambiguity.

Previous methods [4, 44, 45] showed that a suitable representation of particularly traffic scenes is the decomposition into 3D planar surface elements, each one assigned to a rigid body. A rigid body is either the background or a potentially moving object. Following this model, Mono-SF jointly estimates the 3D geometry of each plane and 6D motion of each rigid body considering a) the multi-view geometry by warping the reference image into the consecutive image, b) probabilistic single-view depth estimates, and c) scene model smoothness priors (see Fig. 1). Additionally, an instance segmentation is exploited to detect the set of potentially moving objects.

As an additional contribution, we propose *ProbDepthNet*, a convolutional neural network (CNN) that estimates pixel-wise probability depth distributions from a single image rather than just single depth values such as [10, 14, 19]. Whereas the problem of overconfident estimates is a well-known problem in classification [21], it is typically ignored in probabilistic approaches for regression [17, 31, 32, 34]. Therefore, we propose a novel *recalibration technique*: CalibNet, a small subsequent part of ProbDepthNet, is trained on a hold-out split of the training data to compensate for overfitting effects and to provide well-calibrated distributions.

Our Mono-SF approach is evaluated with respect to several state-of-the-art monocular baselines and an ablation study confirms the importance of the individual components of the proposed optimization framework. Furthermore, ProbDepthNet is validated to provide well-calibrated depth distributions. Our experiments show that several previous probabilistic approaches suffer from overconfident estimates – an effect that could be compensated by adding our proposed CalibNet for recalibration. The suitability of ProbDepthNet for integrating single-view depth information in Mono-SF is confirmed, especially due to the importance of providing single-view depth information in a probabilistic and well-calibrated form.

2. Related work

The works related to the approach presented here are divided into three categories: In the first category are the stereo-based scene flow methods which inspired our Mono-SF scene model and optimization framework. The second category provides an overview of methods for monocular scene reconstruction comprising the baseline methods. Finally, the category of probabilistic deep learning represents works related to the probabilistic design of ProbDepthNet.

Stereo Scene Flow: Scene flow estimation was introduced by Vedula et al. [59, 60] as a joint optimization of 3D geometry and motion of the scene based on a sequence of stereo images. Mostly variational approaches were used subsequently to extend the scene flow concept [3,25,29,50,58,64,65]. However, Vogel et al. [61] were the first that significantly outperformed individual stereo and optical flow methods on their respective tasks for dynamic traffic scenes. They represented the dynamic scene as a collection of rigid moving planar surface elements and jointly optimized the geometry and the motion of each plane considering scene model priors. Menze et al. [44] formulated the problem by a set of rigid moving objects and jointly optimized their motion with the geometry of each plane. This representation is particularly beneficial if the association of planes to objects is supported by an instance segmentation as proposed in [4]. Our Mono-SF model corresponds to these approaches, called *object* [44] or *instance scene flow* [4], but Mono-SF uses only monocular images.

Monocular Scene Reconstruction: Traditionally, monocular scene reconstruction is based on the structure from motion (SfM) principle. The SfM-based approaches can be divided into several categories: First, rigid SfMbased methods estimate the 3D geometry of a rigid scene based on its relative motion to the camera, e.g. a static scene and a moving camera [11, 13, 47, 48, 57]. Second, the nonrigid SfM principle is typically used to derive the deformation of a single object [7, 16, 20]. Third, multi-body SfM is the concept of reconstructing individual moving parts of the scene separately [36,51]. However, the absolute and relative scales of the reconstructions are unknown in general. Scene model assumptions are needed to solve this scale ambiguity, e.g. that moving objects are in contact with the surrounding environment [6,8,51] or that the scene follows a smoothness prior regarding surface and motion [36, 46, 66].

Even though the idea of single-view depth estimation is by far not new [27, 40, 52], the real breakthrough was achieved by usage of deep learning methods. Pioneering, Eigen et al. [10] proposed a CNN that is trained in a supervised manner and estimates the depth in a coarse to fine scheme. Afterward, various self-supervised and unsupervised approaches were proposed using either an image reconstruction loss in a stereo setup [15, 19] or in a monocular image sequence [42, 62, 75, 76]. Fu et al. [14] formulated the depth estimation as an ordinal regression problem, which led to the currently leading approach in the KITTI depth prediction benchmark as reported by [56]. Multitask CNNs that estimate optical flow alongside the depth were proposed [55, 70, 73, 77]. Thereby, both tasks benefit from each other by a combined training loss. DeMoN [57] could also exploit multi-view information for depth estimation during inference. However, it is focused and applied only to static scenes as it just estimates a single camera motion for the whole scene.

Whereas single-view depth estimation and multi-view geometry are mostly taken as individual tasks, a few works combine both. The single-view depth estimation can be useful for scale estimation in monocular visual odometry [2,69,71] or fused with SfM-based depth estimates in static environments [12, 54, 71]. Kumar et al. [37] used singleview depth estimation for depth initialization in a multibody or non-rigid SfM-based approach similar to [36]. Brickwedde et al. [5] proposed a fusion of single-view depth estimates and optical flow to provide a column-wise segmentation in stick-like rigid elements of particularly traffic scenes. In contrast to these methods, Mono-SF is formulated as a scene flow estimation problem and integrates probabilistic single-view depth distributions instead of single depth values.

Probabilistic Deep Learning: The methods of singleview depth estimation mentioned in the previous section do not provide an uncertainty measure or probabilistic distribution of the depth estimates. Kendall and Gal [32] distinguished two kind of uncertainties, epistemic and aleatoric uncertainty. Epistemic uncertainty corresponds to the uncertainty of the model parameters or the ignorance which model generates the training data, whereas aleatoric uncertainty refers to noise in the input data [32]. Malinin et al. [43] extended this definition by introducing the *distribu*tional uncertainty to represent out-of-distribution data. To estimate the extent of aleatoric uncertainty in a regression problem, different strategies have been proposed. First, a probability distribution can be learned by minimizing the negative log-likelihood on the training data [32, 34]. Second, Ilg et al. [31] proposed a single network that is pushed to estimate a complementary set of hypotheses. Thereby, the aleatoric uncertainty is encoded by the empirical distribution of these hypotheses. Third, Gast and Roth [17] replaced each layer with a probabilistic layer to propagate an input uncertainty through the network. The ProbDepthNet method presented here falls under the category of estimating the aleatoric uncertainty with a single network and single inference such as [17, 31, 32, 34]. For classification problems, Guo et al. [21] showed that modern neural networks tend to overfit on the training data, which results in highly overconfident estimates. Recalibration techniques were proposed to compensate for this effect [21, 35, 49].

3. Method

The monocular scene flow estimation method, *Mono-SF*, is designed to combine multi-view geometry with probabilistic single-view depth information in a probabilistic optimization framework. First, a CNN, called *ProbDepthNet*, providing single-view depth information in a probabilistic and well-calibrated form is described. Second, the Mono-SF model and optimization framework are presented.

3.1. Probabilistic Single-View Depth Estimation

To integrate the single-view depth estimates in Mono-SF in a statistical manner, ProbDepthNet is designed to represent the uncertainty of each estimate. Thus, the main objective of ProbDepthNet is not to provide a single depth es-



Figure 2. Overview of ProbDepthNet for probabilistic single-view depth estimation. The architecture consists of two parts: DepthNet and CalibNet for recalibration (blue). Both parts provide a parametrized form (μ_i , s_i / \tilde{s}_i and $\lambda_i / \tilde{\lambda}_i$) of a mixture of Gaussians. Each part is trained on a different split of the training data using a negative log-likelihood loss (orange). The ground truth data is provided by a stereo SGM [26]-based completion of a lidar point cloud (green).

timate, but to provide a probability density function of the depth for each pixel **p** given an input image *I*. The depth is encoded by its inverse form $d = Z^{-1}$, where *Z* is the z-coordinate of the 3D-position in camera coordinates. Prob-DepthNet estimates a pixel-wise probability density function $p_{\mathbf{p}}(d \mid I)$ parameterized as a *mixture of Gaussians*:

$$p_{\mathbf{p}}(d \mid I) = \sum_{i=1}^{K} \lambda_i \cdot \mathcal{N} \left(d - \mu_i, \sigma_i \right)$$
(1)

K represents the number of components, λ_i are the weights, μ_i are the mean values, and σ_i are the variances of the *i*-th component. Compared to a single Gaussian distribution, a mixture model is able to capture more general distributions, e.g. a multimodal distribution. But, the mixture of Gaussians is more an exemplary choice and other parameterizations of a probability distribution can be used as well.

Fig. 2 gives an overview of the architecture, training process and ground truth generation. ProbDepthNet consists of two parts: DepthNet and CalibNet. DepthNet is a fully convolutional ResNet-50 [24] with skip connections between corresponding encoder and decoder layers. The outputs of DepthNet are the parameters of the mixture of Gaussians, whereby the variance is provided in the log-space $s_i = \log \sigma_i$. Additionally, the variances s_i and weights λ_i of DepthNet are recalibrated by CalibNet, which outputs the corresponding recalibrated values \tilde{s}_i and λ_i . CalibNet just consists of five 1×1 convolutional layers: One layer without non-linear activation function to provide a scaled version of the inputs and a residual path with four layers including exponential linear units as activation functions. The number of features of all layers is equal to the number of inputs 2K. Both networks are trained on different splits of the training data to avoid overfitting of DepthNet on the calibration split. The negative log-likelihood loss \mathcal{L} is minimized

during training similar to [32, 34]:

$$\mathcal{L} = \sum_{u,v \in \Omega_{GT}} \left[-\log\left(\sum_{i=1}^{K} \lambda_i \mathcal{N} \left(d_{GT} - \mu_i, \sigma_i \right) \right) \right] \quad (2)$$

 $u, v \in \Omega_{GT}$ are all pixels in the image with valid ground truth depth values d_{GT} and $\mu_i, \lambda_i, \sigma_i$ are the outputs of the trained network.

To overcome the limitations of lidar data in terms of density, range, and field of view, an intermediate fusion based on stereo images is used for ground truth depth generation. First, the lidar point cloud is projected to the image and inconsistent measurements are removed to handle occlusion problems. Second, these sparse depth maps are completed considering a photometric distance between the two stereo images by using an SGM-based approach [26].

ProbDepthNet learns to estimate a pixel-wise depth distribution by observing the depth distribution during the training process. Thereby, the depth distribution captures the aleatoric uncertainty regarding the theory of Kendall and Gal [32]. The aleatoric uncertainty is considered to be the most dominant uncertainty in many vision applications [32]. Our experiments show that CalibNet for recalibration is also applicable to different probabilistic approaches similar to [17,31].

3.2. Monocular Scene Flow

This section presents the Mono-SF optimization framework, structured as follows: First, the decomposition of the scene into piecewise planar surface elements and rigid bodies is described. Second, the optimization is formulated as an energy minimization problem combining a) multi-view geometry-based photometric distance, b) the probabilistic single-view depth estimates of ProbDepthNet and c) scene model smoothness priors. Finally, the inference and initialization of the optimization problem are presented.

Monocular Scene Flow Model: Following previous object scene flow approaches [4, 44, 45], the main assumption is that, in particular, a traffic scene can be approximated by a set of piecewise planar surface elements to represent the structure of the scene and a set of rigid bodies to represent the motion (see Fig. 3). Formally, the reference image is divided into a set of superpixels each one representing a 3D plane. Each 3D plane is defined by its normal $\mathbf{n}_i \in \mathbb{R}^3$, scaled by the inverse distance of the plane to the camera to encode the 3D position \mathbf{X} of each point on the plane by $\mathbf{n}_i^T \mathbf{X} = 1$. The set of rigid bodies consists of the background as well as other traffic participants such as pedestrians or vehicles detected by an instance segmentation. Even though a pedestrian does not undergo a rigid body motion, at a certain scale, it can be approximated by its dominant rigid body transformation as motivated by [45]. Each rigid body is represented by its 6D motion $\mathbf{T}_j \in SE(3)$. Addi-



Figure 3. Variables of Mono-SF model and energy minimization problem are the 6D rigid body motions T_j of moving objects (colored in the left image) and the background as well as the 3D scaled normals n_i of superpixel planes (boundaries in the right image).

tionally, each superpixel is associated with one rigid body and with the pixels \mathcal{R}_i of the corresponding superpixel.

Energy Minimization Problem: The main idea of Mono-SF is that the scene geometry and motion should be consistent in terms of warping the reference image I_0 in the consecutive image I_1 and consistent to the depth distributions $p(d | I_0)$ and $p(d | I_1)$ provided by ProbDepthNet. Formally, Mono-SF jointly optimizes the 6D motion of each rigid body \mathbf{T}_j and 3D normal of each plane \mathbf{n}_i as an energy minimization problem. The energy term E consists of unary data terms $\Phi(\mathbf{p}_0, \mathbf{n}_i, \mathbf{T}_j)$ for each pixel \mathbf{p}_0 and pairwise smoothness terms $\Psi(\mathbf{n}_i, \mathbf{n}_j)$ for each two planes \mathbf{n}_k and \mathbf{n}_l adjacent in the image $k, l \in \mathcal{N}$:

$$E = \sum_{\mathbf{n}_i} \sum_{\mathbf{p}_0 \in \mathcal{R}_i} \Phi(\mathbf{p}_0, \mathbf{n}_i, \mathbf{T}_j) + \sum_{k,l \in \mathcal{N}} \Psi(\mathbf{n}_k, \mathbf{n}_l) \quad (3)$$

 \mathbf{T}_{j} is the rigid body corresponding to the plane \mathbf{n}_{i} .

The unary terms $\Phi(\mathbf{p}_0, \mathbf{n}_i, \mathbf{T}_j)$ consist of two parts. First, $\Phi^{pho}(\mathbf{p}_0, \mathbf{n}_i, \mathbf{T}_j)$ minimizes an appearance-based photometric distance between pixel \mathbf{p}_0 and its projected position in the consecutive image. Second, $\Phi_t^{svd}(\mathbf{p}_0, \mathbf{n}_i, \mathbf{T}_j)$ prefers a 3D position consistent to the estimated depth probabilities of ProbDepthNet at time t = 0 and t = 1:

$$\Phi(\mathbf{p}_{0}, \mathbf{n}_{i}, \mathbf{T}_{j}) = \Theta_{0} \quad \Phi^{pho}(\mathbf{p}_{0}, \mathbf{n}_{i}, \mathbf{T}_{j}) \\ + \Theta_{1} \sum_{t \in \{0, 1\}} \Phi^{svd}_{t}(\mathbf{p}_{0}, \mathbf{n}_{i}, \mathbf{T}_{j}) \quad (4)$$

mh a í

The terms are weighted by Θ_0 or Θ_1 , respectively. The photometric distance $\Phi^{pho}(\mathbf{p}_0, \mathbf{n}_i, \mathbf{T}_j)$ rates the similarity of the two corresponding image positions \mathbf{p}_0 and \mathbf{p}_1 as the hamming distance of their respective 5×5 Census descriptors [74] truncated at τ_0 . The corresponding image coordinates \mathbf{p}_1 in the second image I_1 are defined by a homography [22] considering the 3D normal \mathbf{n}_i and the motion of the corresponding rigid body \mathbf{T}_j :

$$\mathbf{p}_1 = \mathbf{K} (\mathbf{R}_j - \mathbf{t}_j \mathbf{n}_i^T) \mathbf{K}^{-1} \mathbf{p}_0$$
 (5)

 \mathbf{R}_j and \mathbf{t}_j is the decomposition of \mathbf{T}_j into rotation matrix and translation vector. **K** is the intrinsic camera matrix.

The term $\Phi_t^{svd}(\mathbf{p}_0, \mathbf{n}_i, \mathbf{T}_j)$ rates the consistency of the depth of pixel \mathbf{p}_0 based on the ProbDepthNet estimates.

Whereas the depth $d_0(\mathbf{p}_0, \mathbf{n}_i)$ at time t = 0 is directly defined by the corresponding scaled normal vector \mathbf{n}_i , the motion of the corresponding rigid body \mathbf{T}_j needs to be considered to derive the depth $d_1(\mathbf{p}_0, \mathbf{n}_i, \mathbf{T}_j)$ at time t = 1. Both depth values are rated by the negative log-likelihood of the probability provided by ProbDepthNet for their respective image I_t and image coordinate \mathbf{p}_t :

$$\Phi_t^{svd}(\mathbf{p}_0, \mathbf{n}_i, \mathbf{T}_j) = -\log p_{\mathbf{p}_t} \left(d_t(\mathbf{p}_0, \mathbf{n}_i, \mathbf{T}_j) \mid I_t \right) \quad (6)$$

The image coordinates p_1 are again defined as in Eq. (5).

The previous data terms include the single-view depth information and multi-view geometry-based photometric distance. Additionally, scene model priors are integrated similar to [44] as pairwise smoothness terms $\Psi(\mathbf{n}_k, \mathbf{n}_l)$ preferring a smooth structure in terms of depth $\Psi^d(\mathbf{n}_k, \mathbf{n}_l)$ and orientation $\Psi^{ori}(\mathbf{n}_k, \mathbf{n}_l)$, each part weighted by Θ_2 or Θ_3 :

$$\Psi(\mathbf{n}_k, \mathbf{n}_l) = \Theta_2 \Psi^d(\mathbf{n}_k, \mathbf{n}_l) + \Theta_3 \Psi^{ori}(\mathbf{n}_k, \mathbf{n}_l) \quad (7)$$

For each shared boundary pixel $\mathbf{p}_0 \in \mathcal{B}_{k,l}$ of plane \mathbf{n}_k and \mathbf{n}_l , a difference in depth is penalized:

$$\Psi^{d}(\mathbf{n}_{k},\mathbf{n}_{l}) = \sum_{\mathbf{p}_{0}\in\mathcal{B}_{k,l}} \min\left(|d_{0}(\mathbf{p}_{0},\mathbf{n}_{k}) - d_{0}(\mathbf{p}_{0},\mathbf{n}_{l})|,\tau_{1}\right)$$
(8)

Analogously, a smooth orientation of planes adjacent in the image is preferred by measuring the similarity of the normal vectors \mathbf{n}_k and \mathbf{n}_l :

$$\Psi^{ori}(\mathbf{n}_k, \mathbf{n}_l) = \min\left(1 - \frac{|\mathbf{n}_k \mathbf{n}_l|}{||\mathbf{n}_k|||\mathbf{n}_l||}, \tau_2\right)$$
(9)

Both smoothness terms are truncated by τ_1 or τ_2 to regard discontinuities in the depth or orientation, for example between different objects. The hyper-parameters Θ and τ are defined differently according to the rigid body type, background or object, and differently for adjacent planes belonging to different rigid bodies. These dependencies are neglected in the previous equations for ease of reading.

Inference: The scene flow estimation is formulated as the energy minimization problem in Eq. (3). Assuming a suitable initialization, that will be discussed in the next section, an iterative optimization approach can be applied. Following the proposed optimization of the object scene flow methods [4, 44], particle max-product belief propagation is used for 10 iterations with 5 particles for each 6D rigid body motion and 10 particles for each 3D normal vector.

Initialization: The optimization problem needs a suitable initialization of all variables. In the first step, the set of rigid bodies is initialized including their scale-aware 6D motions. Traditionally, the known camera height or an additional inertial measurement unit is used for scale-aware monocular visual odometry in the automotive domain. However, this only provides scale information for

the camera ego-motion. The key idea applied here is to integrate single-view depth information to provide the metric scale. In contrast to [2, 69, 71], we apply this idea additionally for scale-aware pose estimation of moving objects. First, object instances in the images I_0 and I_1 detected by a Mask R-CNN [23] (implementation of [63]) are paired based on sparse flow correspondences ($\mathbf{p}_0^i, \mathbf{p}_1^i$) [18] using a simple voting scheme. Each object instance, as well as the background, builds a rigid body. Second, the 6D motion $\mathbf{T}_j \in SE(3)$ of each rigid body is optimized jointly with a set of 3D points $\mathbf{X}_i \in \mathcal{X}$ (one for each flow correspondence lying in the corresponding instance masks) by minimizing

$$\sum_{\mathbf{X}_i \in \mathcal{X}} \sum_{t \in \{0,1\}} \Theta_4 \Phi_t^{proj}(\mathbf{p}_t^i, \mathbf{X}_i, \mathbf{T}_j) + \Phi_t^{svd}(\mathbf{X}_i, \mathbf{T}_j).$$
(10)

 $\Phi_t^{proj}(\mathbf{p}_t^i, \mathbf{X}_i, \mathbf{T}_j)$ is the reprojection error of \mathbf{X}_i with respect to the flow-based image positions \mathbf{p}_t^i weighted by Θ_4 . $\Phi_t^{svd}(\mathbf{X}_i, \mathbf{T}_j)$ rates the consistency of the 3D points \mathbf{X}_i to the ProbDepthNet estimates analogously to Eq. (6). The energy term of Eq. (10) is optimized using the Levenberg-Marquardt solver implemented in [38].

Subsequently, the set of 3D planes is initialized. First, a dense depth map is computed based on a semi-global matching adapted to the monocular case similarly to [1,67]. Again, the depth estimates are additionally rated by the ProbDepthNet estimates. Second, the superpixels including their 3D normal n_i are initialized using the approach in [68]. The pixels of a plane are enforced to be of the same instance to get a unique association with a rigid body.

4. Experiments

In the first part of this section ProbDepthNet is analyzed: Qualitative results of ProbDepthNet are shown, the generalization capabilities to other datasets are presented and an ablation study confirms the importance of the recalibration technique to provide well-calibrated distributions. In the second part, the Mono-SF optimization framework is evaluated by showing qualitative results and a quantitative evaluation with respect to several state-of-the-art methods. Additionally, two ablation studies confirm the claimed Prob-DepthNet design for Mono-SF and support the importance of the individual components of Mono-SF.

4.1. Probabilistic Single-View Depth Estimation

The experiments are conducted on a ProbDepthNet model trained for the KITTI scene flow training set [45]. The model is trained on 33 sequences of the KITTI raw dataset that are not part of the scene flow set. Around 75% / 25% of the sequences are used for training DepthNet / CalibNet. It is trained for 15 epochs using Adam optimizer [33] with a learning rate of 10^{-4} halved every 5 epochs and a small batch size of 4. The input images are scaled to a size



Figure 4. Exemplary estimates of ProbDepthNet on KITTI scene flow set [44] for the first component of the mixture of Gaussians excluding the weight. The color encodes the inverse depth from close (red) to far (blue) or high variance (red) to low variance (blue).



Figure 5. Generalization of ProbDepthNet (trained on KITTI) on Cityscapes [9] (top) and central crop of Make3D [53] (bottom). The figure shows the estimates based on the left image in the form of the mean depth values μ_0 (middle) and recalibrated log-variances \tilde{s}_0 (right) of the first component.

of 512×256 and a mixture of Gaussians with 8 components is used.

Fig. 4 shows exemplary the output of ProbDepthNet. The variances visually correlate with challenging parts of the scene such as object boundaries or poles. The estimated recalibrated variances \tilde{s}_0 provided by CalibNet are significantly higher than the variances s_0 . The generalization capabilities of ProbDepthNet trained for KITTI are visualized by the qualitative results on the Make3D [53] and Cityscapes [9] dataset in Fig. 5. Please see the supplementary material for more qualitative results and discussions.

The following ablation study analyzes the proposed recalibration by adding the CalibNet trained on a hold-out split. Our proposed training by minimizing the negative loglikelihood (NLL) is related to the approach in [32]. But, to provide a comparison of different probabilistic approaches, the DepthNet part is also trained using a multi-hypothesis strategy ('Hypo [31]') similar to [31] or transformed to its 'assumed density filtering'-counterpart ('ADF [17]') as proposed by [17]. Fig. 6 shows the mean NLL on the KITTI scene flow set (which is not part of the training data) every 1000 training steps. In the bottom plot of Fig. 6, the calibration of the final models is evaluated. The frequency of ground truth depth values inside a given interval should be the same as the cumulative probability of the estimated distribution. The impact of overfitting effects varies among the different approaches - but all approaches suffer from such an effect and provide overconfident estimates. Furthermore, CalibNet is validated as an useful recalibration technique applicable to different probabilistic approaches.

For integration in Mono-SF, a model is additionally pretrained on Cityscapes [9]. Compared to previous nonprobabilistic methods for single-view depth estimation such as [14, 19, 39], the main benefit of ProbDepthNet is providing well-calibrated depth distributions. However, in addition to correct uncertainties, the underlying estimates should have sufficient quality as well. A quantitative evaluation (see supplementary material) shows that the accuracy of the depth estimates represented by the total means of the distributions is comparative to [19, 39] and slightly below [14].

4.2. Monocular Scene Flow

Mono-SF estimates the 3D scene flow from monocular images focusing on dynamic traffic scenes, which means providing the 3D position and 3D motion of each pixel. The following results and evaluations are based on the equiva-



Figure 6. Top: Mean negative log-likelihood (NLL) of ProbDepth-Net on the KITTI scene flow set over the training process; Bottom: Calibration plot comparing the frequency of ground truth depth values lying in a given confidence interval. This frequency is equal to the confidence interval for a perfect calibrated model (dotted line). By including CalibNet for recalibration the overfitting effect is compensated and a better calibrated model achieved.



Figure 7. Exemplary qualitative results of monocular scene flow estimation methods on the KITTI scene flow training set [44]. The top row of each scenario shows the depth values at time t = 0 (left half) and t = 1 (right half) colored from close (red) to far (dark blue). The optical flow is visualized in the bottom row of each scenario. The ground truth is interpolated for visualization purposes.

Method	MRE	D1		D2		Fl		SF					
		bg	fg	all									
GeoNet [73]	20.08	47.03	63.41	49.54	56.25	68.82	58.17	32.43	67.69	37.83	67.69	91.41	71.32
DF-Net [77]	18.95	44.43	57.94	46.50	61.55	61.47	61.54	25.66	37.45	27.47	71.63	82.52	73.30
EveryPixel [70]	-	23.62	27.38	26.81	-	-	-	25.34	28.00	25.74	-	-	-
MirrorFlow [30] + LRC [19]	9.06	25.33	19.83	24.49	35.83	26.15	34.34	9.40	14.22	10.14	40.55	35.17	39.73
$HD^{3}-F^{\dagger}$ [72] + DORN ^{\dagger} [14]	11.18	17.02	37.54	20.16	30.08	40.47	31.67	4.01	6.76	4.43	32.57	46.89	34.76
DMDE [51]	14.6	-	-	-	-	-	-	-	-	-	-	-	-
S. Soup [36]	12.68	-	-	-	-	-	-	-	-	-	-	-	-
MFA [37]	11.82	-	-	-	-	-	-	-	-	-	-	-	-
Mono-Stixels [5]	8.04	18.28	22.06	18.86	22.00	31.19	23.41	9.84	14.36	10.54	24.03	39.13	26.34
Mono-SF (ours)	8.14	15.64	22.72	16.72	17.93	24.71	18.97	12.20	9.90	11.85	20.19	29.40	21.60

MRE: mean relative depth error at t=0 (capped at 50m); D1 and D2: disparity errors at t=0,1; Fl: optical flow errors; SF: scene flow errors

fg: foreground (moving); bg: background (static); all: bg + fg; \dagger : parts of dataset used for training (disregarded for ranking); errors are in percent Table 1. Quantitative evaluation of monocular scene flow methods on the KITTI scene flow training set [45]. The methods are divided into four groups: First, multi-task CNNs; second, combining optical flow and single-view depth estimation as individual tasks; third, multi-body or non-rigid SfM-based approaches; fourth, fusing single-view depth information with multi-view geometry.

lent representation as the depth of each pixel at both times (t = 0, t = 1) and the optical flow. Thereby, the 3D position and the ability of the approaches to predict a 3D point from t = 0 to t = 1 based in its 3D motion is evaluated. Exemplary qualitative results of Mono-SF are shown for the KITTI [45] (see Fig. 7) and Cityscapes dataset [9] (see Fig. 8). Please see the supplementary material for further results.

The quantitative evaluation is based on the KITTI scene flow dataset [45], which reports the frequencies of errors for the depth at time t = 0 (D1) and t = 1 (D2) and the optical flow (Fl). An estimate is considered as an error if it exceeds



Figure 8. Exemplary qualitative result of Mono-SF on a crop of Cityscapes (removing car hood); left: first input image, middle: estimated depth values at time t = 0 (left half) and t = 1 (right half), right: estimated optical flow

a threshold of 3 pixels and 5% in terms of stereo disparity or optical flow endpoint error. Furthermore, an estimate is only defined as a valid scene flow estimate (SF) if it fulfills all the D1, D2, and Fl metrics. All metrics are evaluated separately for moving objects (fg), the static scene (bg) and both combined (all).

We propose four categories of state-of-the-art monocular baseline methods. In the first category are the multi-task networks, GeoNet [73], DF-Net [77] and EveryPixel [70]. These CNNs are trained in an unsupervised manner and are able to provide single-view depth estimates for both images and optical flow estimates. For the GeoNet and DF-Net, their published code and models are used. The results of the EveryPixel approach are stated in their paper [70] (D2 metric is excluded as it seems to be inconsistent). As a second category, single-view depth estimation ('LRC [19]' or 'DORN [14]') and optical flow estimation ('MirrorFlow [30]' or 'HD³-F [72]') are combined as individual tasks. Due to the fact that the published models of 'DORN [14]'

Method	D1-all	D2-all	Fl-all	SF-all
UberATG-DSSF [41]	2.55	4.04	4.73	6.31
ISF [4]	4.46	5.95	6.22	8.08
SGM [26] + SF [28]	6.84	15.60	21.67	24.98
Mono-SF	16.32	19.59	12.77	23.08

Table 2. Results of Mono-SF on the KITTI scene flow test set compared to some stereo-based scene flow estimation methods.

and 'HD³-F [72]' used parts of the dataset for training, these methods are disregarded for ranking. The third group comprises the multi-body or non-rigid SfM-based methods DMDE [51] and S.Soup [36]. The fourth category consists of the methods MFA [37], Mono-Stixel [5] and our Mono-SF approach, which are methods that fuse single-view depth information with multi-view geometry. DMDE, S.Soup, and MFA were only evaluated on its depth estimates capped at 50m using a mean absolute relative error (MRE). For the Mono-Stixel approach, the authors provide us the results on a scene flow metric using MirrorFlow [30] and LRC [19] as inputs. The results of the quantitative evaluation are shown in Table 1. To the best of our knowledge, it is the first time that these methods are evaluated and compared as a scene flow estimation problem. The results show that the methods of the fourth group that combine single-view depth and multi-view geometry outperforms the other methods. Mono-SF shows the best rating on most of the metrics and especially outperforms previous methods on the scene flow (SF) metrics. The approach and implementation of Mono-SF is currently not focused on runtime and needs around 41 seconds per image on a single CPU-core. Mono-SF was also submitted to the KITTI scene flow benchmark (see Table 2). Mono-SF is the first monocular method and would have been ranked at the 13th place with respect to the 21 published stereo scene flow methods.

4.3. Ablation Studies

To analyze the importance of the proposed ProbDepth-Net design, the results of four Mono-SF variants based on different single-view depth estimations are provided in Ta-

Method	D1-all	D2-all	Fl-all	SF-all
Mono-SF (LRC [19])	22.36	26.29	15.10	30.96
Mono-SF (w/o prob. depth)	25.49	28.80	15.04	33.59
Mono-SF (w/o recalib.)	20.32	23.37	15.50	26.91
Mono-SF	16.72	18.97	11.85	21.60

Table 3. Ablation study on ProbDepthNet for Mono-SF. For integrating single-view depth information, ProbDepthNet is more suitable than LRC for single-view depth estimation (improvement over "(LRC [19])"); especially due to the importance of providing single-view depth estimates in a probabilistic (improvement over "(w/o prob. depth)") and well-calibrated form (improvement over "(w/o recalib.)") for Mono-SF.

Energy terms			Results						
Φ^{pho}	Φ^{svd}	Ψ	D1-all	D2-all	Fl-all	SF-all			
-	-	-	18.72	21.30	15.18	25.92			
\checkmark	-	-	21.20	23.41	13.85	26.11			
\checkmark	\checkmark	-	18.65	21.10	13.31	23.67			
\checkmark	\checkmark	\checkmark	16.72	18.97	11.85	21.60			

Table 4. Ablation study on Mono-SF approach. Using the Mono-SF optimization improves the scene flow estimation compared to its initialization (denoted by the row without checkmark). Each term of the energy minimization problem (photometric distance(Φ^{pho}), single-view depth (Φ^{svd}) and smoothness prior (Ψ)) contributes to the final performance.

ble 3. The two Mono-SF variants "Mono-SF (LRC [19])" and "Mono-SF (w/o prob. depth)" utilized CNNs that provide only single-view depth values instead of depth distributions. Whereas "Mono-SF (LRC [19])" is based on the LRC method for single-view depth estimation, "Mono-SF (w/o prob. depth)" is based on the non-probabilistic estimates of ProbDepthNet represented by the total means of the distributions. The depth values are integrated by assuming the same Gaussian distribution (determined on a test set) for all pixels. Mono-SF based on the probabilistic ProbDepthNet ("Mono-SF") outperforms both. This supports the claimed ProbDepthNet design to provide single-view depth estimates in a probabilistic form. Furthermore, the improvements compared to a variant based on the ProbDepthNet excluding CalibNet "Mono-SF (w/o recalib.)" support that the recalibration technique is an essential component.

In Table 4, the individual components of the Mono-SF optimization framework are analyzed by removing some parts of the proposed energy minimization problem (setting their weights to zero). The initialization of Mono-SF described in Sec. 3.2 is denoted by the row without checkmarks. Compared to this initialization, the scene flow formulation of Mono-SF results in further improvement. Additionally, the ablation study shows that each part of the energy term contributes to the final performance; the multiview geometry, the single-view depth information and the scene model smoothness priors.

5. Conclusion

In this paper, we proposed Mono-SF for joint estimation of the 3D geometry and motion of particularly traffic scenes by combining multi-view geometry with single-view depth information. For a sensible statistical integration, we showed the importance of providing single-view depth information in a probabilistic and well-calibrated form, which is made possible by our proposed ProbDepthNet including a novel recalibration technique.

References

- Min Bai, Wenjie Luo, Kaustav Kundu, and Raquel Urtasun. Exploiting semantic information and deep matching for optical flow. In *Proc. of European Conference on Computer Vision (ECCV)*, pages 154–170. Springer, 2016. 5
- [2] Dan Barnes, Will Maddern, Geoffrey Pascoe, and Ingmar Posner. Driven to distraction: Self-supervised distractor learning for robust monocular visual odometry in urban environments. In *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, pages 1894–1900. IEEE, 2018. 2, 5
- [3] Tali Basha, Yael Moses, and Nahum Kiryati. Multi-view scene flow estimation: A view centered variational approach. *International Journal of Computer Vision*, 101(1):6– 21, 2013. 2
- [4] Aseem Behl, Omid Hosseini Jafari, Siva Karthik Mustikovela, Hassan Abu Alhaija, Carsten Rother, and Andreas Geiger. Bounding Boxes, Segmentations and Object Coordinates: How Important is Recognition for 3D Scene Flow Estimation in Autonomous Driving Scenarios? In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2574–2583, 2017. 1, 2, 4, 5, 8
- [5] Fabian Brickwedde, Steffen Abraham, and Rudolf Mester. Exploiting Single Image Depth Prediction for Mono-Stixel Estimation. In Proc. of European Conference of Computer Vision Workshops (ECCV Workshops). IEEE, 2018. 3, 7, 8
- [6] Fabian Brickwedde, Steffen Abraham, and Rudolf Mester. Mono-Stixels: Monocular Depth Reconstruction of Dynamic Street Scenes. In *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–7. IEEE, 2018. 1, 2
- [7] Neil Brikbeck, Dana Cobzas, and Martin Jägersand. Depth and scene flow from a single moving camera. In Proc. of International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT), 2010. 2
- [8] Sebastian Bullinger, Christoph Bodensteiner, Michael Arens, and Rainer Stiefelhagen. 3D Vehicle Trajectory Reconstruction in Monocular Video Data Using Environment Structure Constraints. In Proc. of European Conference on Computer Vision (ECCV), pages 35–50, 2018. 1, 2
- [9] 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 Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 3213–3223, 2016. 6, 7
- [10] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Proc. of Advances in neural information processing systems (NeurIPS)*, pages 2366–2374, 2014. 1, 2
- [11] Jakob Engel, Thomas Schöps, and Daniel Cremers. LSD-SLAM: Large-scale direct monocular SLAM. In Proc. of European Conference on Computer Vision (ECCV), pages 834–849. Springer, 2014. 2
- [12] Fácil, José M and Concha, Alejo and Montesano, Luis and Civera, Javier. Single-View and Multi-View Depth Fusion.

IEEE Robotics and Automation Letters, 2(4):1994–2001, Oct 2017. 1, 3

- [13] Nolang Fanani, Alina Stürck, Matthias Ochs, Henry Bradler, and Rudolf Mester. Predictive monocular odometry (PMO): What is possible without RANSAC and multiframe bundle adjustment? *Image and Vision Computing*, 2017. 2
- [14] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep Ordinal Regression Network for Monocular Depth Estimation. In Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2002–2011, 2018. 1, 2, 6, 7
- [15] Ravi Garg, Gustavo Carneiro, and Ian Reid. Unsupervised CNN for single view depth estimation: Geometry to the rescue. In *Proc. of European Conference on Computer Vision* (ECCV), pages 740–756. Springer, 2016. 2
- [16] Ravi Garg, Anastasios Roussos, and Lourdes Agapito. Dense variational reconstruction of non-rigid surfaces from monocular video. In *Proc. of IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), pages 1272–1279, 2013. 2
- [17] Jochen Gast and Stefan Roth. Lightweight Probabilistic Deep Networks. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 3369–3378, 2018. 2, 3, 4, 6
- [18] Andreas Geiger, Julius Ziegler, and Christoph Stiller. Stereoscan: Dense 3D reconstruction in real-time. In Proc. of IEEE Intelligent Vehicles Symposium (IV), pages 963–968, 2011. 5
- [19] Clement Godard, Oisin Mac Aodha, and Gabriel J. Brostow. Unsupervised Monocular Depth Estimation With Left-Right Consistency. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), July 2017. 1, 2, 6, 7, 8
- [20] Vladislav Golyanik, Aman S Mathur, and Didier Stricker. NRSfM-Flow: Recovering Non-Rigid Scene Flow from Monocular Image Sequences. In Proc. of British Machine Vision Conference (BMVC), 2016. 2
- [21] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On Calibration of Modern Neural Networks. In *Proc. of International Conference on Machine Learning (ICML)*, pages 1321–1330, 2017. 2, 3
- [22] Richard Hartley and Andrew Zisserman. Multiple view geometry in computer vision. Cambridge university press, 2003. 4
- [23] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In Proc. of IEEE International Conference on Computer Vision (ICCV), pages 2980–2988. IEEE, 2017. 5
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 770–778, 2016. 3
- [25] Evan Herbst, Xiaofeng Ren, and Dieter Fox. RGB-D flow: Dense 3-D motion estimation using color and depth. In *Proc.* of *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2276–2282, May 2013. 2
- [26] Heiko Hirschmuller. Accurate and efficient stereo processing by semi-global matching and mutual information. In Proc.

of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 807–814, 2005. 3, 4, 8

- [27] Derek Hoiem, Alexei A Efros, and Martial Hebert. Automatic photo pop-up. In *Proc. of ACM transactions on graphics (TOG)*, volume 24, pages 577–584. ACM, 2005. 2
- [28] Michael Hornacek, Andrew Fitzgibbon, and Carsten Rother. SphereFlow: 6 DoF scene flow from RGB-D pairs. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 3526–3533, 2014.
- [29] Frédéric Huguet and Frédéric Devernay. A variational method for scene flow estimation from stereo sequences. In *Proc. of IEEE International Conference on Computer Vision*, pages 1–7. IEEE, 2007. 2
- [30] Junhwa Hur and Stefan Roth. MirrorFlow: Exploiting symmetries in joint optical flow and occlusion estimation. In Proc. of International Conference on Computer Vision (ICCV), 2017. 7, 8
- [31] Eddy Ilg, Ozgun Cicek, Silvio Galesso, Aaron Klein, Osama Makansi, Frank Hutter, and Thomas Brox. Uncertainty Estimates and Multi-Hypotheses Networks for Optical Flow. In *Proc. of European Conference on Computer Vision (ECCV)*, 2018. 2, 3, 4, 6
- [32] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Proc. of Advances in neural information processing systems* (*NeurIPS*), pages 5574–5584, 2017. 2, 3, 4, 6
- [33] Diederik P Kingma and Jimmy Lei Ba. Adam: Amethod for stochastic optimization. In *Proc. of International Conference for Learning Representations (ICLR)*, 2014. 5
- [34] Maria Klodt and Andrea Vedaldi. Supervising the new with the old: learning SFM from SFM. In *Proc. of European Conference on Computer Vision (ECCV)*, pages 698–713, 2018. 2, 3, 4
- [35] Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. Accurate Uncertainties for Deep Learning Using Calibrated Regression. In Proc. of International Conference on Machine Learning (ICML), pages 2801–2809, 2018. 3
- [36] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Monocular dense 3D reconstruction of a complex dynamic scene from two perspective frames. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, pages 4649–4657, 2017. 1, 2, 3, 7, 8
- [37] Suryansh Kumar, Ram Srivatsav Ghorakavi, Yuchao Dai, and Hongdong Li. A Motion Free Approach to Dense Depth Estimation in Complex Dynamic Scene. arXiv preprint arXiv:1902.03791, 2019. 3, 7, 8
- [38] Rainer Kümmerle, Giorgio Grisetti, Hauke Strasdat, Kurt Konolige, and Wolfram Burgard. g 2 o: A general framework for graph optimization. In *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, pages 3607– 3613. IEEE, 2011. 5
- [39] Yevhen Kuznietsov, Jörg Stückler, and Bastian Leibe. Semisupervised deep learning for monocular depth map prediction. In Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 6647–6655, 2017. 6
- [40] Lubor Ladicky, Jianbo Shi, and Marc Pollefeys. Pulling things out of perspective. In Proceedings of the IEEE Confer-

ence on Computer Vision and Pattern Recognition (CVPR), pages 89–96, 2014. 2

- [41] Wei-Chiu Ma, Shenlong Wang, Rui Hu, Yuwen Xiong, and Raquel Urtasun. Deep rigid instance scene flow. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2019. 8
- [42] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised Learning of Depth and Ego-Motion from Monocular Video Using 3D Geometric Constraints. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 5667–5675, 2018. 2
- [43] Andrey Malinin and Mark Gales. Predictive uncertainty estimation via prior networks. In Proc. of Advances in Neural Information Processing Systems (NeurIPS), pages 7047–7058, 2018. 3
- [44] Moritz Menze and Andreas Geiger. Object Scene Flow for Autonomous Vehicles. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015. 1, 2, 4, 5, 6, 7
- [45] Moritz Menze, Christian Heipke, and Andreas Geiger. Object Scene Flow. *ISPRS Journal of Photogrammetry and Remote Sensing*, 140:60 – 76, 2018. Geospatial Computer Vision. 2, 4, 5, 7
- [46] Amar Mitiche, Yosra Mathlouthi, and Ismail Ben Ayed. Monocular Concurrent Recovery of Structure and Motion Scene Flow. *Frontiers in ICT*, 2:16, 2015. 1, 2
- [47] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. ORB-SLAM: a versatile and accurate monocular SLAM system. *IEEE Transactions on Robotics*, 31(5):1147– 1163, 2015. 2
- [48] Fabio Irigon Pereira, Gustavo Ilha, Joel Luft, Marcelo Negreiros, and Altamiro Susin. Monocular Visual Odometry with Cyclic Estimation. In *Graphics, Patterns and Images* (SIBGRAPI), 2017 30th SIBGRAPI Conference on, pages 1– 6. IEEE, 2017. 2
- [49] John Platt. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. Advances in large margin classifiers, 10(3):61–74, 1999. 3
- [50] Jean-Philippe Pons, Renaud Keriven, and Olivier Faugeras. Multi-view stereo reconstruction and scene flow estimation with a global image-based matching score. *International Journal of Computer Vision*, 72(2):179–193, 2007. 2
- [51] René Ranftl, Vibhav Vineet, Qifeng Chen, and Vladlen Koltun. Dense monocular depth estimation in complex dynamic scenes. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4058–4066, 2016. 1, 2, 7, 8
- [52] Ashutosh Saxena, Sung H Chung, and Andrew Y Ng. Learning depth from single monocular images. In Proc. of Advances in Neural Information Processing Systems (NeurIPS), volume 18, pages 1–8, 2005. 2
- [53] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Make3D: Learning 3D Scene Structure from a Single Still Image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):824–840, May 2009. 6
- [54] Keisuke Tateno, Federico Tombari, Iro Laina, and Nassir Navab. CNN-SLAM: Real-time dense monocular SLAM

with learned depth prediction. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, 2017. 1, 3

- [55] Qianru Teng, Yimin Chen, and Chen Huang. Occlusion-Aware Unsupervised Learning of Monocular Depth, Optical Flow and Camera Pose with Geometric Constraints. *Future Internet*, 10(10):92, 2018. 2
- [56] Jonas Uhrig, N Schneider, L Schneider, U Franke, Thomas Brox, and A Geiger. Sparsity Invariant CNNs. In Proc. of IEEE International Conference on 3D Vision (3DV), 2017. 2
- [57] 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 Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), July 2017. 2
- [58] Levi Valgaerts, Andrés Bruhn, Henning Zimmer, Joachim Weickert, Carsten Stoll, and Christian Theobalt. Joint estimation of motion, structure and geometry from stereo sequences. In *Proc. of European Conference on Computer Vision (ECCV)*, pages 568–581. Springer, 2010. 2
- [59] Sundar Vedula, Simon Baker, Peter Rander, Robert Collins, and Takeo Kanade. Three-dimensional scene flow. In *Proc. of IEEE International Conference on Computer Vision* (*ICCV*), volume 2, pages 722–729. IEEE, 1999. 1, 2
- [60] Sundar Vedula, Peter Rander, Robert Collins, and Takeo Kanade. Three-dimensional scene flow. *IEEE transactions* on pattern analysis and machine intelligence, 27(3):475– 480, 2005. 1, 2
- [61] Christoph Vogel, Konrad Schindler, and Stefan Roth. Piecewise rigid scene flow. In Proc. of the IEEE International Conference on Computer Vision (ICCV), pages 1377–1384, 2013. 1, 2
- [62] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning Depth from Monocular Videos using Direct Methods. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2022–2030, 2018. 2
- [63] Naiyan Wang. An MXNet implementation of Mask R-CNN. https://github.com/TuSimple/mx-maskrcnn, 2018. [accessed June, 25 2018]. 5
- [64] Andreas Wedel, Thomas Brox, Tobi Vaudrey, Clemens Rabe, Uwe Franke, and Daniel Cremers. Stereoscopic scene flow computation for 3D motion understanding. *International Journal of Computer Vision*, 95(1):29–51, 2011. 2
- [65] Andreas Wedel, Clemens Rabe, Tobi Vaudrey, Thomas Brox, Uwe Franke, and Daniel Cremers. Efficient dense scene flow from sparse or dense stereo data. In Proc. of European Conference on Computer Vision (ECCV), pages 739– 751. Springer, 2008. 2
- [66] Degui Xiao, Qiuwei Yang, Bing Yang, and Wei Wei. Monocular scene flow estimation via variational method. *Multimedia Tools and Applications*, 76(8):10575–10597, Apr 2017. 1, 2
- [67] Koichiro Yamaguchi, David McAllester, and Raquel Urtasun. Robust monocular epipolar flow estimation. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 1862–1869, 2013. 5

- [68] Koichiro Yamaguchi, David McAllester, and Raquel Urtasun. Efficient Joint Segmentation, Occlusion Labeling, Stereo and Flow Estimation. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, Proc. of European Conference on Computer Vision (ECCV), pages 756– 771, Cham, 2014. Springer International Publishing. 5
- [69] Nan Yang, Rui Wang, Jörg Stückler, and Daniel Cremers. Deep virtual stereo odometry: Leveraging deep depth prediction for monocular direct sparse odometry. In Proc. of European Conference on Computer Vision (ECCV), pages 835–852. Springer, 2018. 2, 5
- [70] Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, and Ram Nevatia. Every Pixel Counts: Unsupervised Geometry Learning with Holistic 3D Motion Understanding. In Proc. of European Conference on Computer Vision Workshops (ECCV Workshops), 2018. 2, 7
- [71] Xiaochuan Yin, Xiangwei Wang, Xiaoguo Du, and Qijun Chen. Scale recovery for monocular visual odometry using depth estimated with deep convolutional neural fields. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5870–5878, 2017. 1, 2, 3, 5
- [72] Zhichao Yin, Trevor Darrell, and Fisher Yu. Hierarchical discrete distribution decomposition for match density estimation. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2019. 7, 8
- [73] Zhichao Yin and Jianping Shi. GeoNet: Unsupervised Learning of Dense Depth, Optical Flow and Camera Pose. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), volume 2, 2018. 2, 7
- [74] Ramin Zabih and John Woodfill. Non-parametric local transforms for computing visual correspondence. In *Proc. of European Conference on Computer Vision (ECCV)*, pages 151– 158. Springer, 1994. 4
- [75] 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 *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 340–349, 2018. 2
- [76] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised Learning of Depth and Ego-Motion from Video. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017. 2
- [77] Yuliang Zou, Zelun Luo, and Jia-Bin Huang. DF-Net: Unsupervised Joint Learning of Depth and Flow using Cross-Network Consistency. In Proc. of European Conference on Computer Vision (ECCV), pages 36–53, 2018. 2, 7