

Self-supervised Visual-LiDAR Odometry with Flip Consistency

Bin Li¹, Mu Hu¹, Shuling Wang¹, Lianghao Wang², and Xiaojin Gong^{*†1}

¹College of Information Science and Electronic Engineering, Zhejiang University, China,

²Zhejiang Lab, Hangzhou, China

¹{3130102392, muhu, 11831041, gongxj}@zju.edu.cn, ²wang_sunsy@163.com;

Abstract

Most learning-based methods estimate ego-motion by utilizing visual sensors, which suffer from dramatic lighting variations and textureless scenarios. In this paper, we incorporate sparse but accurate depth measurements obtained from lidars to overcome the limitation of visual methods. To this end, we design a self-supervised visual-lidar odometry (Self-VLO) framework. It takes both monocular images and sparse depth maps projected from 3D lidar points as input, and produces pose and depth estimations in an end-to-end learning manner, without using any ground truth labels. To effectively fuse two modalities, we design a two-pathway encoder to extract features from visual and depth images and fuse the encoded features with those in decoders at multiple scales by our fusion module. We also adopt a siamese architecture and design an adaptively weighted flip consistency loss to facilitate the self-supervised learning of our VLO. Experiments on the KITTI odometry benchmark show that the proposed approach outperforms all self-supervised visual or lidar odometries. It also performs better than fully supervised VOs, demonstrating the power of fusion.

1. Introduction

Ego-motion or pose estimation is a key problem in simultaneous localization and mapping (SLAM), which plays an important role in various applications such as autonomous driving[36], 3D reconstruction [7], and augmented reality [17]. Most traditional methods [30, 36, 51, 38] address this problem by following a feature extraction and matching pipeline, relying on geometry models for optimization. These methods have shown excellent perfor-

mance, but they depend on hand-crafted modules and lack the ability of adapting to all challenging scenarios. Recently, learning-based techniques [53, 46, 20, 23, 3] have attracted more and more research interest. Although these methods perform inferior to traditional counterparts, they demonstrate advantages in dealing with tough cases such as rotation-only motion or tracking lost scenarios.

To date, most learning-based methods use visual sensors. These methods, also referred to as deep visual odometries (VOs) [53, 21, 20], take advantage of rich information in monocular or stereo images, but are sensitive to lighting conditions. Several works [19, 23, 3] adopt ranging sensors such as lidars for ego-motion estimation. These sensors can provide accurate depth measurements and are robust to the change of lighting conditions, but their data are sparse. Therefore, visual and ranging sensors are highly complementary so that it is desirable to integrate them to achieve better performance. In addition, considering that most autonomous vehicles are equipped with both visual and ranging sensors nowadays, the study of visual-lidar odometries (VLOs) will be of great value for academic research and industrial applications.

In this work, we propose a learning-based method to implement a visual-lidar odometry. In order to effectively fuse visual and lidar information, we first project 3D lidar points to obtain image-aligned sparse depth maps. Then, we design a network composed of a two-pathway encoder to extract features from a monocular image sequence and a sparse depth map sequence, respectively. The extracted features are fused at multiple scales and employed to jointly predict pose and dense depth values. Considering the difficulty of obtaining ground-truth labels, we opt to learn our VLO under self-supervision. To this end, we adopt a siamese network architecture and design a flip consistency loss to facilitate the self-supervised learning.

The main contributions of our work are summarized as follows:

- To the best of our knowledge, this work is the first end-

^{*}The corresponding author

[†]This work was supported in part by the Natural Science Foundation of Zhejiang Province, China under Grant No. LY17F010007 and LY18F010004.

to-end learning-based visual-lidar odometry. It inputs monocular images and sparse depth maps, and outputs both depth and pose estimations.

- We adopt a siamese network architecture and design a flip consistency loss to learn our VLO under self-supervision, without using any ground-truth pose and depth labels.
- Evaluations on the KITTI odometry benchmark [6] show that our proposed method outperforms all self-supervised visual or lidar odometries. It also performs better than all fully supervised VOs, demonstrating the power of fusion.

2. Related Work

2.1. Visual & Lidar Odometry

Visual odometry aims to estimate ego-motion based on images. Traditional VOs can be classified into feature-based or direct methods. The former [30, 17, 27] follows a standard pipeline including feature extraction and matching, motion estimation, and local/global optimization. The direct methods [29, 14] minimize a photometric error that takes all pixels into account, assuming that scenes are stationary and of a Lambertian reflectance. However, these assumptions are hard to be fully satisfied so that most direct methods perform inferior to their feature-based counterparts.

Recently, deep learning techniques have been firstly applied to supervised VO [39, 40, 43, 44] that requires ground-truth pose labels. Considering the hardness of obtaining ground-truth, more and more works focus on unsupervised VO. The pioneer work, SfMLearner [53], takes a monocular video sequence as input to jointly learn depth and pose networks by minimizing a photometric loss. Following it, various end-to-end learning methods have been developed. For instance, UnDeepVO[34], Zhu et al.[54], and Zhan et al.[50] take stereo sequences as training inputs to address the scale ambiguity in monocular VO. SAVO [21] and Li et al. [20] adopt recurrent neural networks (RNNs) to encode temporal information. Additional predictions such as optical flow [46], disparity map [54], and motion mask [34, 9] are also included to address motion or occlusion problem. In addition, some recent works such as DVSO [28] and D3VO [45] integrate the end-to-end pose estimation with traditional global optimization modules to boost VO performance further.

Lidar odometry estimates the relative pose between two 3D point cloud frames. Although a lidar provides accurate depth information, the irregular sparsity of point clouds makes LO a challenging problem. Most previous lidar odometry works, such as LOAM [51], LeGO-LOAM [38], and LOL [35], are based on traditional pipelines and achieve

state-of-the-art performance. Recently, deep learning techniques have also been applied to this task. For instance, L³-Net [23] replaces each module in the traditional pipeline with a deep neural network in order to learn from data. DeepVCP [22], DeepLO [3] and LO-Net [19] design various end-to-end learning frameworks. The former [22] adopts a 3D convolutional neural network (CNN) for feature extraction and matching, while the latter two [3, 19] project point clouds into 2D maps via a cylindrical or spherical projection and use 2D CNNs for processing. Most learning-based LOs are performed under full supervision, except DeepLO [3].

Visual-lidar odometry takes advantage of visual and lidar information for ego-motion estimation. Up to now, most of the existing VLOs are based on traditional methods. For instance, V-LOAM [52] integrates the results of VO and LO, both of which adopt the feature-based pipelines, to boost the performance. LIMO [10] provides a feature-based monocular VO with the scale estimated from lidar data to achieve ego-motion estimation. Recently, there are several visual and lidar based methods making use of learning techniques. Yue et al. [47] exploit a deep network to enrich 3D point clouds with the guidance of high-resolution images. But they employ a traditional point cloud registration method NDT [26] to estimate pose. Vitor et al. [11] incorporates depth supervision into a self-supervised VO to produce scale-aware pose estimation. But the input of their network only contains visual images. In this work, we propose a self-supervised visual-lidar odometry. It takes a sequence of monocular images and sparse depth maps projected from corresponding 3D lidar points as inputs to learn a depth and pose estimation network in an end-to-end manner, without using any ground-truth pose and depth labels.

2.2. Depth Estimation

Depth estimation works can be classified into depth prediction and depth completion categories, according to the KITTI benchmark [6]. The former [5] estimates depth from a single color image, while the latter takes a color image and a sparse depth map as the inputs to produce a dense depth map. In contrast to learning-based VOs [39, 53, 50] that input an image sequence, most depth estimation works take a single frame as input. They focus more on the design of network architectures to encode multi-scale contexts [18], semantics [4] or surface normals [33] for performance improvement. Moreover, most depth estimation methods are based on supervised learning, requiring ground-truth depth maps that are extremely hard to obtain. Recently, unsupervised methods have also been developed. For instance, Ma et al. [25] implements self-supervised depth completion via also taking a sequence of images and sparse depth maps as inputs. Their work shares the same inputs with ours, but it uses a traditional method to estimate ego-motion between

two frames.

2.3. Consistency Regularization

Consistency regularization aims to improve a model's performance via enforcing the consistency between two output predictions. For instance, some stereo matching [49] and depth estimation [31, 8, 24] works take a stereo pair as inputs and place a left-right consistency on their corresponding outputs. Optical flow [55] takes consecutive images as inputs and impose a forward-backward consistency. Recently, the consistency between the outputs of an original image and its transformed (e.g. rescaled, rotated, or flipped) version has been exploited to construct supervision for various self-supervised learning tasks in image classification [2, 12], object detection [15], semantic segmentation [41], etc. In this work, we adopt the flip consistency and design an adaptively weighted consistency loss specifically for self-supervised visual-lidar odometry.

3. The Proposed Method

In this section, we first present the network architecture designed for our visual-lidar odometry. Then we introduce the proposed losses to achieve the self-supervised learning.

3.1. Network Architecture

Visual-lidar fusion: When designing the network architecture, the first problem we need to address is how to fuse visual and lidar modalities. Inspired by recent depth completion works [33, 37], we adopt an encoder-decoder fusion scheme and design the network as shown in Figure 1. First, in contrast to deep LOs [3, 19] that take a cylindrical or spherical projection, we project each 3D point cloud into its aligned image frame and obtain a 2D sparse depth map. Although this projection manner leaves out the 3D points that are not in the camera's field of view, it facilitates the fusion of color and depth information.

We refer to our visual-lidar odometry network as VLONet. It has a two-pathway encoder. Each takes either an image sequence or a depth map sequence as input. The features extracted in each pathway are fused by a point-wise addition for the prediction of relative pose between two consecutive frames. In addition, the depth and visual features are also fused together with decoder features via a feature fusion module at each scale, as shown in Figure 1, to predict a dense depth map. In contrast to most deep VOs [53, 34, 50] that construct two separate networks for depth and pose prediction, our depth and pose estimation share the same encoder, which can not only reduce the computational cost but also extract more powerful features.

A siamese architecture for the self-supervised learning: Inspired by recent unsupervised learning works [2, 12, 15, 41], we adopt a siamese architecture for our self-supervised VLO. More specifically, the architecture con-

sists of two siamese VLONets, as shown in Figure 2. One network takes original image and depth sequences as input, and the other takes the horizontally flipped images and depth maps as input. These two networks share the same structure and parameters. Based on this architecture, we can impose a consistency regularization to facilitate our self-supervised learning. This siamese architecture is used for training, while only one VLONet is needed for test.

3.2. Depth and Pose Prediction Losses

We first introduce the losses for the prediction of depth and pose in a single VLONet. As no ground-truth labels are available, we adopt a view synthesis loss \mathcal{L}_{vs} , together with a depth fidelity loss \mathcal{L}_{df} and a depth smoothness loss \mathcal{L}_{ds} for supervision. Therefore, the entire loss of a single VLONet is defined by

$$\mathcal{L}_{vlo} = \lambda_{vs}\mathcal{L}_{vs} + \lambda_{df}\mathcal{L}_{df} + \lambda_{ds}\mathcal{L}_{ds}, \quad (1)$$

in which λ_{vs} , λ_{df} , and λ_{ds} are scaling factors to balance the terms. For the purpose of self-containedness, we briefly introduce each term as follows.

View synthesis loss, which is extensively used in unsupervised VOs [53, 34], measures the difference between a target image and an image synthesized from a source image. Assuming that each sequence contains $2N + 1$ frames, we denote the input image sequence by $\{I_{t-N}, \dots, I_{t-1}, I_t, I_{t+1}, \dots, I_{t+N}\}$ and the depth map sequence by $\{D_{t-N}, \dots, D_{t-1}, D_t, D_{t+1}, \dots, D_{t+N}\}$. The frame I_t and D_t are the target views and the rest are the source views denoted by I_s and D_s , $s \in \{t-N, \dots, t+N\}$, $s \neq t$. The VLONet predicts a dense depth map \hat{D}_t and a relative pose $\hat{P}_{t \rightarrow s} = [\hat{t}_x, \hat{t}_y, \hat{t}_z, \hat{r}_x, \hat{r}_y, \hat{r}_z]$ for each target and source pair. Here, $[\hat{t}_x, \hat{t}_y, \hat{t}_z]$ is a 3D translational vector and $[\hat{r}_x, \hat{r}_y, \hat{r}_z]$ is a 3D Euler angle, from which we can construct a 4×4 transformation matrix $\hat{T}_{t \rightarrow s}$.

Given a pixel p_t in the target image I_t , we can get its corresponding pixel p_s in a source image I_s by the following transform:

$$p_s = K\hat{T}_{t \rightarrow s}\hat{D}_t(p_t)K^{-1}p_t, \quad (2)$$

where

$$K = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \quad (3)$$

is a camera intrinsic matrix provided in training. Further, by taking the differentiable bilinear sampling mechanism [53], we warp the source image into the target frame to get the synthesized image $\tilde{I}_{s \rightarrow t}$, in which $\tilde{I}_{s \rightarrow t}(p_t) = I_s(\langle p_s \rangle)$.

Then, a view synthesis loss is designed to measure the difference between I_t and $\tilde{I}_{s \rightarrow t}$. To this end, we combine a L1 error with a SSIM [48] and define the loss as follows:

$$\mathcal{L}_{vs} = \sum_s (1 - \alpha_s) \|I_t - \tilde{I}_{s \rightarrow t}\|_1 + \alpha_s (1 - SSIM(I_t, \tilde{I}_{s \rightarrow t})), \quad (4)$$

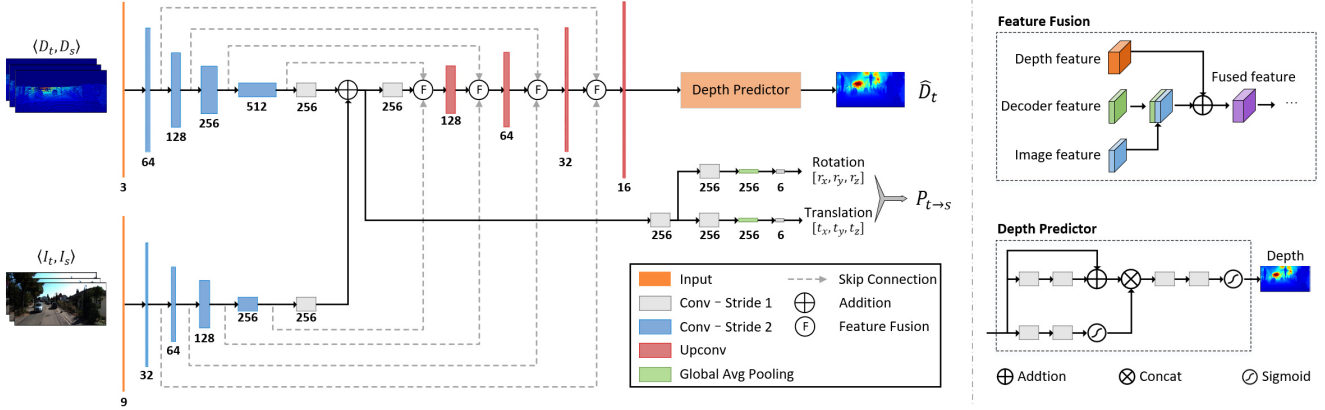


Figure 1. The network architecture of our visual-lidar odometry, referred to as VLONet. It follows an encoder-decoder structure and consists of three major components: a two-pathway encoder, a pose estimator and a depth prediction decoder.

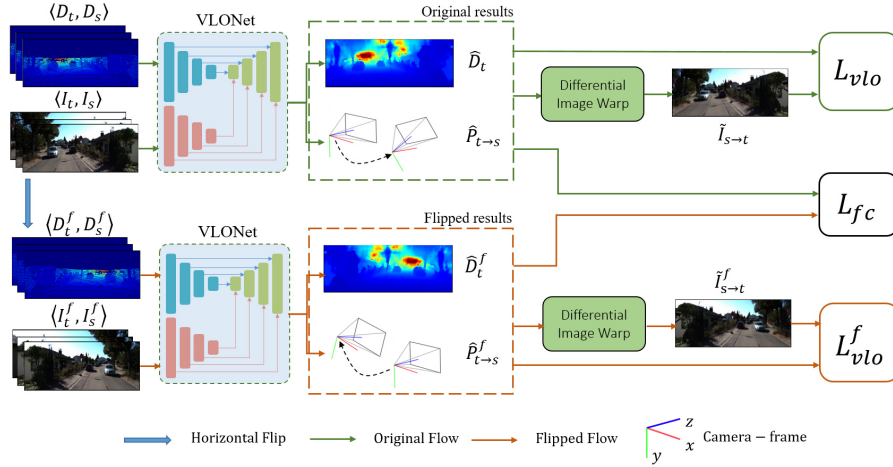


Figure 2. An overview of our Self-VLO. It consists of two VLONets that share the same structures and parameters with each other. One VLONet takes original sequences as input and the other inputs horizontally flipped sequences. Each network is supervised with corresponding depth and pose prediction losses. There is also a flip consistency loss imposed to constrain the outputs of two VLONets.

where α_s is a scalar to balance the two terms. According to Equation (2), we know that this loss provides a supervision for both depth and pose estimations.

Depth fidelity loss is to ensure the predicted dense depth map have the same values with the input sparse depth map at these valid pixels. Therefore, this loss is defined by

$$\mathcal{L}_{ds} = \|M(D_t) \odot (D_t - \hat{D}_t)\|_1, \quad (5)$$

in which $M(D_t) = \mathbb{1}(D_t > 0)$ indicates the valid pixels in the sparse map. Note that the supervision of depth comes from the inputs and no extra ground-truth depth labels are required. Therefore, our method is self-supervised.

Depth smoothness loss is imposed to improve the quality of the predicted dense depth map. Considering the correlation between color and depth images, we introduce an image-guided loss to smooth depth while keep sharp on ob-

ject boundaries. This loss is defined by

$$\mathcal{L}_{sm} = \sum_p \exp(-|\nabla_x I_t(p)| |\nabla_x \hat{D}_t(p)|) \quad (6)$$

$$+ \exp(-|\nabla_y I_t(p)| |\nabla_y \hat{D}_t(p)|), \quad (7)$$

where ∇_x, ∇_y denote the gradients.

3.3. The Flip Consistency Loss

By taking advantage of the siamese architecture, we also impose a consistency regularization to facilitate our self-supervised training. It is based on the following assumption [41]: if there is an affine transformation $A(\cdot)$ on the input, the corresponding output is inclined to be equivariant. That is, $G(A(I, D)) = A(G(I, D))$, where $G(\cdot)$ denotes our VLONet function mapping the input to the output. In this work, we only take the horizontal flip into account for simplicity. Therefore, we first derive the relation between

two outputs and then design a flip consistency loss to regularize the two outputs.

Recall that the original input sequences are $\langle I_t, I_s \rangle$ and $\langle D_t, D_s \rangle$, and their outputs are \hat{D}_t and $\hat{P}_{t \rightarrow s} = [\hat{t}_x, \hat{t}_y, \hat{t}_z, \hat{r}_x, \hat{r}_y, \hat{r}_z]$. We horizontally flip each image and depth frame while keep their temporal order unchanged. The flipped sequences are denoted by $\langle I_t^f, I_s^f \rangle$ and $\langle D_t^f, D_s^f \rangle$, and the corresponding outputs are denoted by \hat{D}_t^f and $\hat{P}_{t \rightarrow s}^f = [\hat{t}_x^f, \hat{t}_y^f, \hat{t}_z^f, \hat{r}_x^f, \hat{r}_y^f, \hat{r}_z^f]$. These two outputs have the following relations:

$$\begin{cases} \hat{D}_t^f = \text{flip}(\hat{D}_t) \\ \hat{P}_{t \rightarrow s}^f = [\hat{t}_x^f, \hat{t}_y^f, \hat{t}_z^f, \hat{r}_x^f, \hat{r}_y^f, \hat{r}_z^f] \\ \quad = [-\hat{t}_x, \hat{t}_y, \hat{t}_z, \hat{r}_x, -\hat{r}_y, -\hat{r}_z] \end{cases} \quad (8)$$

where $\text{flip}(\cdot)$ denotes the horizontal flip operation. Note that, when an image is flipped, the intrinsic matrix of the camera is also changed into the following one

$$K^f = \begin{bmatrix} f_x & 0 & W - c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \quad (9)$$

where W is the image width.

Based on the above-mentioned relations, we place consistency constraints on both depth and pose estimation results. Therefore, the flip consistency loss is defined as follows.

$$\mathcal{L}_{fc} = \mathcal{L}_{dfc} + \mathcal{L}_{pfc}, \quad (10)$$

in which

$$\mathcal{L}_{dfc} = \|\text{flip}(\hat{D}_t) - \hat{D}_t^f\|_1, \quad (11)$$

and

$$\begin{aligned} \mathcal{L}_{pfc} = & \|\hat{t}_x + \hat{t}_x^f + \hat{t}_y - \hat{t}_y^f + \hat{t}_z - \hat{t}_z^f\|_1 \\ & + \alpha_r \|\hat{r}_x - \hat{r}_x^f + \hat{r}_y + \hat{r}_y^f + \hat{r}_z + \hat{r}_z^f\|_1 \end{aligned} \quad (12)$$

are two losses to constrain the consistency of the depth and pose estimations, respectively. α_r is a scalar to balance the translational error and the rotational error.

3.4. The Entire Loss

In summary, when training with the siamese network architecture, we have two depth and pose estimation losses \mathcal{L}_{vlo} and \mathcal{L}_{vlo}^f , respectively, for the original and flipped sequences, together with a flip consistency loss \mathcal{L}_{fc} to constrain the two outputs. Therefore, the entire loss is

$$\mathcal{L} = \mathcal{L}_{vlo} + \mathcal{L}_{vlo}^f + \phi \mathcal{L}_{fc}. \quad (13)$$

Here,

$$\phi = \lambda_{fc} \cdot \exp\left(-\frac{\mathcal{L}_{vs}}{\sigma}\right) \quad (14)$$

is an adaptive weight controlled by the view synthesis loss \mathcal{L}_{vs} to balance the prediction losses and the consistency loss. λ_{fc} and σ are two scalars to ensure the adaptive weight ϕ range from 0 to 1.

The reason to employ this adaptive weighting scheme [42] is because the consistency loss \mathcal{L}_{fc} is prone to result in a trivial solution with all zeros in depth and pose values. Involving the consistency loss prematurely may make the model parameters oscillate with time when the depth and pose estimates are incorrect, leading to a slow convergence at the early epochs. Therefore, we use this adaptive weighting scheme to vary the weight during training time by letting ϕ be inversely proportional to the \mathcal{L}_{vs} . ϕ is small when \mathcal{L}_{vs} is large and ϕ will be close to 1 when \mathcal{L}_{vs} goes to 0 as the training converges.

4. Experiments

4.1. Experiment Setting

Dataset and evaluation metrics. We evaluate our self-supervised visual-lidar odometry and its variants on the KITTI odometry benchmark [6]. As the common practice [53, 34, 54], we take sequence 00-08 for training and test on sequence 09 and 10. To evaluate the pose estimation results, we compute the average translational error (%) and rotational error (deg/100m) on all possible sub-sequences of length (100, 200, \dots , 800) meters, following the official criteria provided in the KITTI benchmark.

Implementation details. We implement the proposed model based on the PyTorch [32] framework. In all experiments, the length of each input sequence is 3. Each image or depth map is resized into 192×624 , which is half of the original resolution, to save computational cost. We train our full model using a single NVIDIA GTX 2080Ti. The full model takes 50 hours or so for training. It infers both depth and pose at a rate of 40Hz during test time.

The details of our network, including the operations and channel size, are marked in Figure 1. In addition, we use batch normalization [13] and ReLU activation for all convolutional layers except those in the prediction layers. During training, we use the Adam optimizer [16] with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and mini-batch size of 4 to train the network for 180K iterations. The initial learning rate starts from 0.0002 and decreases by half for every 70K iterations. We empirically set the hyper-parameters as follows: $\lambda_{vs} = 2.0$, $\lambda_{df} = 0.2$, $\lambda_{ds} = 40.0$, $\alpha_s = 0.85$, and $\alpha_r = 2.0$. In the adaptive weight ϕ , we set the scaling factors $\lambda_{fc} = 2.5$ and $\sigma = 0.2$.

Data augmentation. In our network, each input sequence is a 3-frame snippet. If we only select consecutive frames as inputs, then the low-speed snippets may take the majority while the high-speed snippets take a very small portion, as shown in Figure 3(a). This imbalanced distri-

Models	Inputs	DA	SA	FC	Seq.09		Seq.10		Mean	
					t_{rel}	r_{rel}	t_{rel}	r_{rel}	t_{rel}	r_{rel}
VLO1	M+L				4.33	1.72	3.30	1.40	3.82	1.56
VLO2	M+L	✓			3.42	1.59	4.04	1.40	3.73	1.50
VLO3	M+L	✓	✓		3.21	1.33	3.56	1.14	3.39	1.24
VLO4	M+L	✓	✓	✓	2.58	1.13	2.67	1.28	2.62	1.21
VO1	M				11.34	3.15	16.70	5.04	14.02	4.10
VO2	M	✓			9.83	3.53	14.92	4.43	12.38	3.98
VO3	M	✓	✓		8.30	2.56	15.90	4.92	12.10	3.74
VO4	M	✓	✓	✓	7.00	2.41	11.74	3.33	9.37	2.87

Table 1. Comparison of the proposed method and its variants. Here, M denotes the input taking only monocular image sequences, M+L is the input taking both monocular images and depth maps projected from lidar points. DA denotes the data augmentation strategy, SA is the siamese architecture, and FC is the flip consistency loss. t_{rel} is the translational error (%) and r_{rel} is the rotational error (deg/100m). The best results of VLOs or VO4 are marked in bold.

bution makes a model perform poorly at high-speed scenarios. To address this issue, we opt to enlarge the sampling interval by sampling 3 frames, respectively, at time $t - 2$, t , and $t + 2$ with a probability $p = 0.6$. By this means, the speed distribution of two consecutive frames in 3-frame snippets is more balanced, as shown in Figure 3(b). In addition, we also observe that extremely low-speed snippets (e.g. the speed between two frames is slower than 10km/h) may hurt the performance. The reason is that the geometry model of pose estimation degenerates when two frames are too close to each other. Therefore, we leave out these snippets during training. Except these, no other data augmentation techniques such as random scaling, cropping or flipping are adopted.




Figure 3. The distributions of the speed between two consecutive frames in 3-frame snippets sampled from the KITTI odometry benchmark (Seq. 00-08). DA refers to the data augmentation strategy designed in our work.

4.2. Ablation Studies

We first conduct a series of experiments to investigate the effectiveness of each component proposed in our model. To this end, we test the following four model variants: 1) VLO1: the model using only a single VLONet and without our data augmentation (DA) strategy; 2) VLO2: the model using one VLONet; 3) VLO3: the model using the siamese architecture composed of two VLONets but without the flip consistency loss; 4) VLO4: the full model. All of the latter

three models are trained with DA. In addition, to validate the effectiveness of visual-lidar fusion, we also investigate four visual models, denoted by VO1-4, which are respectively corresponding to VLO1-4 but take only monocular images as input.

Table 1 presents the pose estimation results of all variants. Note that the VO models take only monocular images as input that lead to a scale ambiguity problem, we recover the scales by a post-processing step as done in [53, 34]. When comparing each VLO model with the corresponding VO counterpart, we observe that the fusion of visual and lidar information improves the pose estimation performance dramatically. In addition, all proposed components including the data augmentation strategy, the siamese architecture, and the flip consistency loss, gradually boost the performance in both VLOs and VO4, demonstrating their effectiveness. Figure 4 presents the trajectories of sequence 09 and 10, produced by four model variants. As shown in the figure, both VLO1 and VLO4 generate the trajectories much closer to the ground truth than the VO models, and the full model VLO4 is the closest.




Figure 4. Trajectories of sequence 09 and 10.

In order to take a close look at how the proposed models make improvements, Figure 5 also plots the translational and rotational errors along each axis, averaged on sub-sequences of sequence 09 with a length of (100, 200, ..., 800) meters. We see that the fusion of vi-

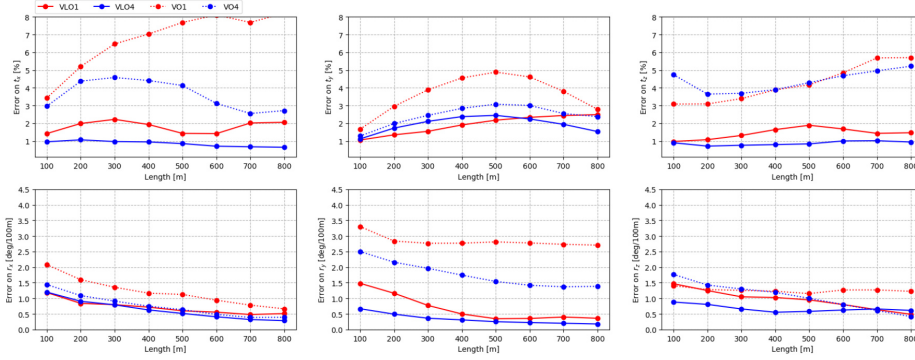


Figure 5. Translational and rotational errors along each axis, averaged on sub-sequences of sequence 09 with different length. Top row: the averaged translational errors along X , Y , Z axis, respectively. Bottom row: the averaged rotational errors along pitch, yaw, and roll.

	Method	Inputs	Supervision	Seq.09		Seq.10		Mean(09-10)	
				t_{rel}	r_{rel}	t_{rel}	r_{rel}	t_{rel}	r_{rel}
Supervised	DeepVO[39]	Mono	Pose	-	-	8.11	8.83	8.11	8.83
	ESP-VO[40]	Mono	Pose	-	-	9.77	10.2	9.77	10.2
	GFS-VO[43]	Mono	Pose	-	-	6.32	2.33	6.32	2.33
	Xue et al.[44]	Mono	Pose	-	-	3.94	1.72	3.94	1.72
	LO-Net[19]	LiDAR	Pose	1.37	0.58	1.80	0.93	1.59	0.76
Self-supervised	SfMLearner[53]	Mono	-	18.77	3.21	14.33	3.30	16.55	3.26
	UnDeepVO[34]	Stereo	-	7.01	3.61	10.63	4.65	8.82	4.13
	Zhu et al.[54]	Stereo	-	4.66	1.69	6.30	1.59	5.48	1.64
	Zhan et al.[50]	Stereo	-	11.92	3.60	12.62	3.43	12.27	3.52
	Gordon et al.[9]	Mono	-	3.10	-	5.40	-	4.25	-
	SAVO [21]	Mono	-	9.52	3.64	6.45	2.41	7.99	3.03
	Li et al. [20]	Mono	-	5.89	3.34	4.79	0.83	5.34	2.09
	DeepLO[3]	LiDAR	-	4.87	1.95	5.02	1.83	4.95	1.89
	Self-VLO	Mono+LiDAR	-	2.58	1.13	2.67	1.28	2.62	1.21

Table 2. Comparison of the proposed method with state-of-the-art learning based methods. Note that the self-supervised methods are trained on sequence 00-08. For the supervised methods, [39, 40, 43, 44] are trained on sequence 00, 02, 08 and 09, and [19] is trained on sequence 00-06.

sual and lidar information, as shown by VLO1 and VLO4 in Figure 5, dramatically improves the performance on t_x (along the lateral direction), t_z (along the driving direction), and r_y (yaw), which are the directions dominate a vehicle’s motion. Besides, the errors of VLO4 are more stable than the other variants, indicating that the full model is more robust to the change of driving speed and road scenarios.

4.3. Comparison with State-of-the-arts

In this section, we compare our proposed full model (named as Self-VLO) with both fully supervised and unsupervised/self-supervised methods. All previous learning-based methods either take visual images or lidar data as input. Therefore, the state-of-the-art fully supervised methods include four visual odometries [39, 40, 43, 44] and one lidar odometry [19]. The self-supervised methods include seven VOs [53, 34, 54, 50, 9, 21, 20] and one LO [3]. These visual methods may take monocular or stereo image sequences as input for training, but all of them infer

pose from monocular images during test.

Table 2 summarizes the comparison results. From it we see that our Self-VLO outperforms all visual based methods, including both fully supervised and self-supervised ones, on test sequences 09 and 10. Moreover, Self-VLO also performs better than the unsupervised LO [3] although it takes 360° field of view into account. These comparisons demonstrate the effectiveness of our method.

5. Conclusions

In this paper, we have presented a self-supervised visual-lidar odometry. It takes advantage of rich visual information and accurate depth information to improve the performance of ego-motion estimation. We adopt a siamese architecture and design an adaptively weighted flip consistency loss to facilitate the learning of our VLO under self-supervision, requiring no manual annotations. Experiments on KITTI validate the effectiveness of our proposed method.

References

- [1] Jens Behley and Cyrill Stachniss. Efficient surfel-based slam using 3d laser range data in urban environments. In *Robotics: Science and Systems*, 2018.
- [2] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [3] Y. Cho, G. Kim, and A. Kim. Unsupervised geometry-aware deep lidar odometry. In *ICRA*, 2020.
- [4] D. Eigen and R. Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, 2015.
- [5] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NIPS*, 2014.
- [6] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012.
- [7] A. Geiger, J. Ziegler, and C. Stiller. Stereoscan: Dense 3d reconstruction in real-time. In *IVS*, 2011.
- [8] C. Godard, O. Mac, and G. Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017.
- [9] A. Gordon, H. Li, R. Jonschkowski, and A. Angelova. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *ICCV*, 2019.
- [10] J. Graeter, A. Wilczynski, and M. Lauer. Limo: Lidar-monocular visual odometry. In *IROS*, 2018.
- [11] V. Guizilini, J. Li, R. Ambrus, S. Pillai, and A. Gaidon. Robust semi-supervised monocular depth estimation with re-projected distances. In *CoRL*, 2019.
- [12] H. Guo, K. Zheng, X. Fan, H. Yu, and S. Wang. Visual attention consistency under image transforms for multi-label image classification. In *CVPR*, 2019.
- [13] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [14] T. Schöps, J. Engel, and D. Cremers. LSD-SLAM: Large-Scale Direct monocular SLAM. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 8690 LNCS, pages 834–849, 2014.
- [15] J. Jeong, S. Lee, J. Kim, and N. Kwak. Consistency-based semi-supervised learning for object detection. In *NIPS*, 2019.
- [16] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [17] G. Klein and D. Murray. Parallel tracking and mapping for small ar workspaces. In *ISMAR*, 2007.
- [18] A. Li, Z. Yuan, Y. Ling, W. Chi, S. Zhang, and C. Zhang. A multi-scale guided cascade hourglass network for depth completion. In *WACV*, 2020.
- [19] Q. Li, S. Chen, C. Wang, X. Li, C. Wen, M. Cheng, and J. Li. Lo-net: Deep real-time lidar odometry. In *CVPR*, 2019.
- [20] S. Li, X. Wang, Y. Cao, F. Xue, Z. Yan, and H. Zha. Self-supervised deep visual odometry with online adaptation. In *CVPR*, 2020.
- [21] S. Li, F. Xue, X. Wang, Z. Yan, and H. Zha. Sequential adversarial learning for self-supervised deep visual odometry. In *ICCV*, 2019.
- [22] W. Lu, G. Wan, Y. Zhou, X. Fu, P. Yuan, and S. Song. Deep-vcp: An end-to-end deep neural network for point cloud registration. In *ICCV*, 2019.
- [23] Y. Lu, W. and Zhou, G. Wan, S. Hou, and S. Song. L3-net: Towards learning based lidar localization for autonomous driving. In *CVPR*, 2019.
- [24] T. Hassner M. Goldman and S. Avidan. Learn stereo, infer mono: Siamese networks for self-supervised, monocular, depth estimation. In *CVPRW*, 2019.
- [25] F. Ma, G. V. Cavalheiro, and S. Karaman. Self-supervised sparse-to-dense: Self-supervised depth completion from lidar and monocular camera. In *ICRA*, 2019.
- [26] M. Magnusson, A. Lilienthal, and T. Duckett. Scan registration for autonomous mining vehicles using 3d-ndt. *Journal of Field Robotics*, 24(10):803–827, 2007.
- [27] R. Mur-Artal and J. Tardos. ORB-SLAM2: An Open-Source SLAM System for Monocular, Stereo, and RGB-D Cameras. *IEEE Transactions on Robotics*, 33(5):1255–1262, 2017.
- [28] J. Stuckler N. Yang, R. Wang and D. Cremers. Deep virtual stereo odometry: Leveraging deep depth prediction for monocular direct sparse odometry. In *ECCV*, pages 817–833, 2018.
- [29] R. Newcombe, S. Lovegrove, and A. Davison. DTAM: Dense tracking and mapping in real-time. In *ICCV*, pages 2320–2327, 2011.
- [30] D. Nister, O. Naroditsky, and J. Bergen. Visual odometry. In *CVPR*, 2004.
- [31] Y. Liu P. Chen, A. Liu and Y. Wang. Towards scene understanding: Unsupervised monocular depth estimation with semantic-aware representation. In *CVPR*, 2019.
- [32] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, Z. Killeen, T. and Lin, N. Gimeshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejjani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeuroIPS*, 2019.
- [33] J. Qiu, Z. Cui, Y. Zhang, X. Zhang, S. Liu, B. Zeng, and M. Pollefeys. Deeplidar: Deep surface normal guided depth prediction for outdoor scene from sparse lidar data and single color image. In *CVPR*, 2019.
- [34] Z. Long R. Li, S. Wang and D. Gu. Undeepvo: Monocular visual odometry through unsupervised deep learning. In *ICRA*, 2018.
- [35] D. Rozenberszki and A. L. Majdik. Lol: Lidar-only odometry and localization in 3d point cloud maps. In *ICRA*, 2020.
- [36] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of rgb-d slam systems. In *IROS*, 2012.
- [37] J. Tang, F. Tian, W. Feng, J. Li, and P. Tan. Learning guided convolutional network for depth completion. 2019.
- [38] T. Shan and B. Englot. Lego-loam: Lightweight and ground-optimized lidar odometry and mapping on variable terrain. In *IROS*, 2018.

- [39] S. Wang, R. Clark, H. Wen, and N. Trigoni. Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks. In *ICRA*, 2017.
- [40] S. Wang, R. Clark, H. Wen, and N. Trigoni. End-to-end, sequence-to-sequence probabilistic visual odometry through deep neural networks. *International Journal of Robotics Research*, 37(4-5):513–542, 2018.
- [41] Y. Wang, J. Zhang, M. Kan, S. Shan, and X. Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *CVPR*, 2020.
- [42] A. Wong and S. Soatto. Bilateral cyclic constraint and adaptive regularization for unsupervised monocular depth prediction. In *CVPR*, 2019.
- [43] F. Xue, Q. Wang, X. Wang, W. Dong, J. Wang, and H. Zha. Guided feature selection for deep visual odometry. In *Asian Conference on Computer Vision*, pages 293–308. Springer, 2018.
- [44] F. Xue, X. Wang, S. Li, Q. Wang, J. Wang, and H. Zha. Beyond tracking: Selecting memory and refining poses for deep visual odometry. In *CVPR*, 2019.
- [45] N. Yang, L. Stumberg, R. Wang, and D. Cremers. D3vo: Deep depth, deep pose and deep uncertainty for monocular visual odometry. In *CVPR*, 2020.
- [46] Z. Yin and J. Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *CVPR*, 2018.
- [47] J. Yue, W. Wen, J. Han, and L. Hsu. Lidar data enrichment using deep learning based on high-resolution image: An approach to achieve high-performance lidar slam using low-cost lidar. *arXiv preprint arXiv:2008.03694*, 2020.
- [48] H. Sheikh Z. Wang, A. Bovik and E. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [49] J. Zbontar and Y. LeCun. Stereo matching by training a convolutional neural network to compare image patches. *Journal of Machine Learning Research*, 17:1–32, 2016.
- [50] H. Zhan, R. Garg, C. Weerasekera, K. Li, H. Agarwal, and I. Reid. Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In *CVPR*, 2018.
- [51] J. Zhang and S. Singh. Loam: Lidar odometry and mapping in real-time. In *RSS*, 2014.
- [52] J. Zhang and S. Singh. Visual-lidar odometry and mapping: Low drift, robust, and fast. In *ICRA*, 2015.
- [53] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017.
- [54] A. Zhu, W. Liu, Z. Wang, V. Kumar, and K. Daniilidis. Robustness meets deep learning: An end-to-end hybrid pipeline for unsupervised learning of egomotion. *arXiv preprint arXiv:1812.08351*, 2018.
- [55] Y. Zou, Z. Luo, and J. Huang. Df-net: Unsupervised joint learning of depth and flow using cross-task consistency. In *ECCV*, 2018.