

Attentional Separation-and-Aggregation Network for Self-supervised Depth-Pose Learning in Dynamic Scenes

Feng Gao¹, Jincheng Yu¹, Hao Shen², Yu Wang¹, Huazhong Yang¹

{gao-f19, yjc16}@mails.tsinghua.edu.cn, shenhao04@meituan.com

{yu-wang, yanghz}@tsinghua.edu.cn

¹ Tsinghua University ² Meituan-Dianping Group

Abstract: Learning depth and ego-motion from unlabeled videos via self-supervision from epipolar projection can improve the robustness and accuracy of the 3D perception and localization of vision-based robots. However, the rigid projection computed by ego-motion cannot represent all scene points, such as points on moving objects, leading to false guidance in these regions. To address this problem, we propose an **Attentional Separation-and-Aggregation Network (ASANet)**, which can learn to distinguish and extract the scene’s static and dynamic characteristics via the attention mechanism. We further propose a novel MotionNet with an ASANet as the encoder, followed by two separate decoders, to estimate the camera’s ego-motion and the scene’s dynamic motion field. Then, we introduce an auto-selecting approach to detect the moving objects for dynamic-aware learning automatically. Empirical experiments demonstrate that our method can achieve the state-of-the-art performance on the KITTI benchmark.

Keywords: Self-Supervised Depth-Pose Learning, Attentional Separation-and-Aggregation Network, Auto-Selecting Mechanism

1 Introduction

Inferring the 3D structure and camera’s ego-motion from videos is significant in 3D computer vision tasks, especially for robotics, autonomous driving, and augmented reality. The traditional methods take the idea of structure-from-motion, using the stereo matching to recover the depth map, and solving the essential matrix by the perspective-n-point (PnP) [1] method. These geometric methods usually rely on the hand-crafted features with correspondence search and epipolar geometry. Recently, learning-based depth estimation [2, 3, 4, 5] and ego-motion estimation [6, 7, 8, 9] methods have emerged and achieved great success. However, the requirement of the high-precision ground-truth limits their flexibility and transportability. To utilize the pervasive unlabeled monocular videos, Zhou et al. [10] propose to learn depth and ego-motion jointly from videos in a self-supervised manner. Along this line, many works have put forward many improvements in consistency constraints [11, 12] and awareness of occlusion [13]. Another remaining challenge is moving objects, which can destroy the scene’s rigid global transformation caused by the camera’s ego-motion.

To tackle the issue of the moving objects, most previous works [14, 15] used a pre-trained segmentation model to preprocess the data, and then separately estimated the movements of the individuals. The pre-trained model, like Mask R-CNN [16], is always bulky and introduces the exterior annotated knowledge, making the possible online adaptation [17] difficult, which is one of the main advantages of self-supervised learning. In this paper, we aim to design a dynamic-aware learning system under the pure self-supervision from monocular videos. Since Gordon et al. [15] have demonstrated that the CNN-based model can estimate the camera’s ego-motion and the motion field of selected regions, monocular videos naturally contain all features required to complete these two estimations. We note that ego-motion estimation, or called visual odometry, relies on static elements in the scene, while dynamic features contribute to motion field estimation for moving objects.

In order to distinguish and exploit the static and dynamic features of the scene, we propose an attentional separation-and-aggregation network (ASANet). There are two weights-sharing feature

propagation paths inside the ASANet, corresponding to the static and dynamic characteristics. In each layer of each path, the features are first decomposed into static and dynamic elements using soft attention, and then the decomposed static and dynamic features from two paths are aggregated into a new pair of features and fed into the subsequent layers. Through multiple separation-and-aggregation operations, ASANet finally outputs pure static and dynamic features. Leveraging the ASANet as the encoder, our extended MotionNet can simultaneously estimate the ego-motion of the camera and the scene’s motion field by two separate decoders. Also, we design an auto-selecting mechanism to select the appropriate transformation to build supervision automatically, i.e., using ego-motion for projecting stationary elements and motion field for transforming the moving objects. We conduct extensive experiments to verify our method’s performance and compare our method with the existing methods, showing that our method can achieve better performance at a lower cost.

Our contributions in this work can be summarized as follows:

- We propose an Attentional Separation-and-Aggregation Network (ASANet), which can extract the static and dynamic features simultaneously from stacked images.
- We propose an auto-selecting approach to eliminate the moving objects issue in a self-supervised manner. To our best knowledge, we are the first to alleviate effect from arbitrary moving objects in self-supervised depth-pose learning from monocular videos without any exterior knowledge.
- Our method achieves state-of-the-art performance on the depth estimation and end-to-end visual odometry tasks of the KITTI benchmark [18, 19] with only a small additional cost.

2 Related Work

2.1 Supervised learning for depth and ego-motion

Supervised methods use precise ground-truth to train neural networks, and these ground-truth data often require expensive instruments for measurement, such as LiDAR used to acquire depth and Inertial Navigation Systems (INS) to record poses. Eigen et al. [2] proposed the first CNN-based model to estimate the dense depth from a single raw image. After that, some subsequent works tried some methods to help the model learn better, such as using conditional random fields (CRFs) to fuse multi-scale information [3], introducing structural attention [4], or treating depth estimation as an ordered regression task [5]. The ego-motion estimation is more famous as visual odometry (VO), which measures the agent’s posture transformation from two views. Since VO is a typical sequence-to-sequence task, the recurrent units are widely used in the existing supervised VO methods. Wang et al. [7] proposed the first end-to-end learning-based VO approach by leveraging recurrent convolutional neural networks (RCNNs). Jiao et al. [20] and Xue et al. [21] improved the model architecture to make better use of temporal geometric relationships. Despite the rapid evolution of these supervised methods, the demand for high-precision ground-truth limits their further development.

2.2 Self-supervised joint depth-pose learning

Self-supervised joint depth-pose learning is an emerging topic in recent years, which links depth and ego-motion via structure-from-motion (SfM). Zhou et al. [10] proposed a pioneering work in this area. On this basis, some of the follow-up works put forward new consistency constraints [11, 12], and some added occlusion processing in the training framework [13, 22]. Moving objects are also a concern in this genre because they are not satisfied with the transformation calculated by ego-motion. The auto-masking approach proposed by Godard et al. [13] can avoid the influence of objects that are relatively stationary with the agent. Casser et al. [14] and Gordon et al. [15] proposed to segment all instances in advance by a pre-trained Mask R-CNN [16] model and then to estimate the motion field for these individuals. As a similar self-supervised video learning task, self-supervised scene flow learning is also facing moving objects, where some works [23, 24] use residual flow to narrow the gap caused by moving objects. Unlike these methods, our method uses the attention mechanism to explore the dynamic features of the scene directly and automatically adapts to the static and dynamic elements during training.

On the other hand, some recent works attempted to incorporate geometric VO into the self-supervised manner. Zhen et al. [25] first proposed to utilize a FlowNet to predict bi-directional

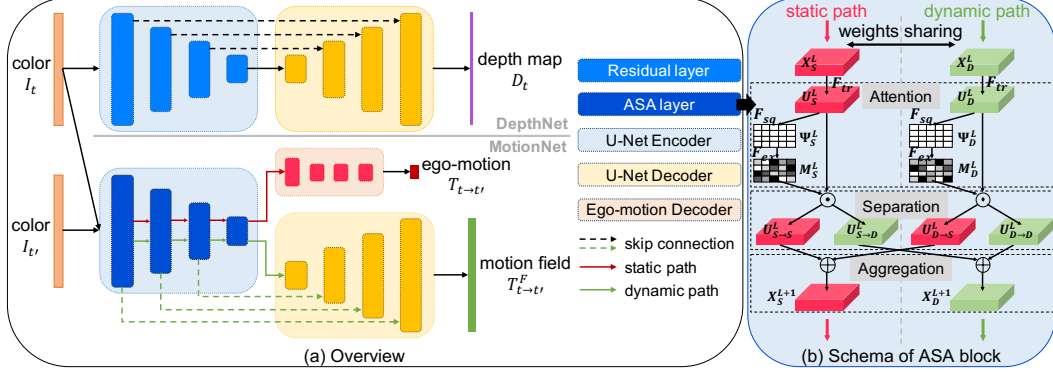


Figure 1: **(a) Overview of the joint learning framework:** a DepthNet in the shape of U-Net and a MotionNet composed of ASANet encoder and two separate decoders. **(b) Schema of the proposed ASA block**, where the red and green volumes represent static and dynamic features, respectively.

optical flows and samples corresponding points from flows with least forward-backward flow inconsistency to solve fundamental matrix. Zhao et al. [22] further introduced the self-supervised manner to this paradigm and enhanced the generalization ability. However, these geometric VO methods all spend most of their time on feature extraction (or flow estimation) and matching, which makes the method difficult to achieve high running speed. Despite a performance gap from geometric VO, end-to-end VO (or called direct VO) has good potential and can run in real-time.

3 Methodology

In this section, we describe our proposed Attentional Separation-and-Aggregation Network (ASANet) and the details of the training process. The overall framework of our joint learning system is depicted in Fig. 1. Following Godard et al. [13], we take the ResNet [26] to encode the input image to leverage the pretraining on ImageNet [27], and then use the U-shaped decoder to reproject the coded feature into the inverse depth, or called disparity. Besides, we extend the commonly used PoseNet to MotionNet, which estimates the ego-motion from static features and estimates the motion field from dynamics.

3.1 Attentional Separation-and-Aggregation Network

As aforementioned, only part of features between two frames contributes to the ego-motion estimation, while the others represent the dynamic characteristics good for handling moving objects. We propose an Attentional Separation-and-Aggregation Network (ASANet) to separate and extract static and dynamic features from adjacent frames. In detail, we design a novel ASA block that uses spatial soft attention in the feature space to weight the static elements and therefore can use the opposite attention to weight the dynamic elements in the meantime.

We diagram the schema of our proposed ASA block in Fig. 1 (b), which contains three main operations: attention generation, feature separation, and aggregation. An ASA block receives two features (static and dynamic) from the last block, notated as X_S^L and X_D^L . Each of the two features generates a propagation path and undergoes the following operations. For any given transformation F_{tr} , one of the input features X_i^L converts to the intermediate feature U_i^L , where $U_i^L \in \mathbb{R}^{H \times W \times C}$, and $i \in \{S, D\}$ indicates which feature path is current. In our implementation, we take the non-identity branch of a residual block as F_{tr} like in SE-ResNet [28]. Then for **(a) attention generation**, we first squeeze the feature volume U_i^L at the channel dimension, denoted as F_{sq} in Fig. 1, to focus on the spatial relationships, thus we can get a spatial correspondence descriptor $\Psi_i^L \in \mathbb{R}^{H \times W \times 1}$. Thereafter, a set of convolutions F_{ex} process the descriptor Ψ_i^L to produce a probability map M_i^L , each element of which expresses how likely it describes the static characteristics. For **(b) feature separation**, we apply the static-likely mask M_i^L and the dynamic-likely mask $(1 - M_i^L)$ to U_i^L , obtaining two new separated features $U_{i \rightarrow S}^L$ and $U_{i \rightarrow D}^L$. At the end of an ASA block, we **(c) aggregate** two new pairs of features from two paths, $X_S^{L+1} = U_{S \rightarrow S}^L + U_{D \rightarrow S}^L$ and $X_D^{L+1} = U_{S \rightarrow D}^L + U_{D \rightarrow D}^L$, and finally output the purer static and dynamic features to the rest of the network.

Using an ASANet as the encoder, our MotionNet feeds the extracted static and dynamic features into two independent decoders to regress the ego-motion and dynamic motion field, respectively, among which the motion field decoder and the ASANet constitute a U-shaped structure. Through the skip connections, the impure dynamic features from the shallow layers provide the necessary structural information and some residual static features for restoring the scene’s motion field.

3.2 Self-supervised objective

In line with SfMLearner [10], we formulate the self-supervised depth-pose learning as novel view synthesis. Let’s denote $\langle I_{t-1}, I_t, I_{t+1} \rangle$ as a training snippet with I_t being the target view and the rest being the source view $I_{t'} (t' \neq t)$. The DepthNet estimates the disparity D_t from a single image I_t and then restores it to the dense depth map by $\hat{D}_t = 1/D_t$, while the MotionNet estimates the ego-motion $T_{t \rightarrow t'}$ and the motion field $T_{t \rightarrow t'}^F$ from the stacked images. Then the estimated depth map D and relative pose \mathcal{T} can warp one view into another according to the epipolar geometry:

$$p_{t'} \sim K\mathcal{T}\hat{D}_t(p_t)K^{-1}p_t \quad (1)$$

where K denotes the camera intrinsics, and $p_t, p_{t'}$ are the homogeneous coordinates of pixels in I_t and $I_{t'}$. \mathcal{T} is a short version of $\mathcal{T}_{t \rightarrow t'}(p_t)$ for the sake of simplicity, equal to either the ego-motion $T_{t \rightarrow t'}$ or one element of the motion field $T_{t \rightarrow t'}^F(p_t)$, simply notated as T and T^F below. Utilizing the view synthesis, we can optimize our models by minimizing the following consistency errors.

Photometric consistency loss measures the difference between the target image I_t and the synthesized view $\hat{I}_t^{t'} = \mathcal{F}(I_{t'}, \hat{D}_t, \mathcal{T})$, being a combination of a L1 and SSIM [29] loss, where \mathcal{F} is the warping operation. Then the photometric error (pe) is defined as

$$pe(I_t, \hat{I}_t^{t'}) = \frac{\alpha}{2}(1 - SSIM(I_t, \hat{I}_t^{t'})) + (1 - \alpha)\|I_t - \hat{I}_t^{t'}\|_1 \quad (2)$$

Inspired by Godard et al. [13], we compute the per-pixel minimum error across all source views instead of averaging them and incorporate the auto-masking term (Eq. (3)), to overcome the issues of out-of-view pixels and occlusions.

$$\mathcal{M}_{auto} = \llbracket \min_{t'} pe(I_t, \hat{I}_t^{t'}) < \min_{t'} pe(I_t, I_{t'}) \rrbracket \quad (3)$$

$$\mathcal{L}_{pe} = \mathcal{M}_{auto} \odot \min_{t'} pe(I_t, \hat{I}_t^{t'}) \quad (4)$$

where $\llbracket \cdot \rrbracket$ indicates the Iverson bracket and \odot is the element-wise product.

Geometric consistency loss is proposed in [30] to ensure a consistent scale between different training snippets. Similar to photometric consistency, the warping operation is applied but on the depth map rather than the RGB image. We can compute the geometric error (ge) by

$$ge(\hat{D}_t, \hat{D}_{t'}, \mathcal{T}) = \frac{|\hat{D}_t^{t'} - \hat{D}_{t'}|}{\hat{D}_t^{t'} + \hat{D}_{t'}} \quad (5)$$

where $\hat{D}_t^{t'} = \mathcal{F}(\hat{D}_t, \mathcal{T})$ is the warped depth map from \hat{D}_t and $\hat{D}_{t'}$ is the interpolated depth map of $I_{t'}$ aligned with the warping flow. Here, we also compute the per-pixel minimum error instead of the average.

$$\mathcal{L}_{ge} = \mathcal{M}_{auto} \odot \min_{t'} ge(\hat{D}_t, \hat{D}_{t'}, \mathcal{T}) \quad (6)$$

Smoothness regularization is commonly used in previous works [10, 13, 23]. We adopt two edge-aware smoothness losses here, where one is a first-order term to regularize the depth maps [10, 13], and the other is a second-order term for the motion field [23].

$$\mathcal{L}_{ds} = |\nabla_x^1 D_t^*| e^{-|\nabla_x^1 I_t|} + |\nabla_y^1 D_t^*| e^{-|\nabla_y^1 I_t|} \quad (7)$$

$$\mathcal{L}_{fs} = |\nabla_x^2 F_{t \rightarrow t'}| e^{-|\nabla_x^2 \hat{D}_t|} + |\nabla_y^2 F_{t \rightarrow t'}| e^{-|\nabla_y^2 \hat{D}_t|} \quad (8)$$

where $D_t^* = D_t / \bar{D}_t$ is the mean-normalized disparity map, $\hat{D}_t = 1/D_t$ is the restored depth map, and $F_{t \rightarrow t'}$ is the warping flow derived from the epipolar projection, an intermediate for smoothing the motion field:

$$F_{t \rightarrow t'}(p_t) = K T_{t \rightarrow t'}^F(p_t) \hat{D}_t(p_t) K^{-1} p_t - p_t \quad (9)$$



Figure 2: **Visual results to show the effect of our dynamic-aware learning¹.** The first two rows are sample images and corresponding estimated depth maps. The third row shows the results of auto-masking \mathcal{M}_{auto} and our proposed auto-selecting $\mathcal{M}_{dynamic}$ approach. The last row illustrates the photometric error maps before and after the dynamic pixel-wise projection is incorporated into the rigid global projection via our auto-selecting mechanism.

3.3 Dynamic-aware learning

Dynamic awareness is critical for view synthesis, because the epipolar projection produced by camera motion is invalid for moving objects. To solve this problem, we propose an auto-selecting mechanism that automatically selects the transformation source for each point by comparing the consistency error from ego-motion and that from the motion field. We can calculate the selective mask according to the following formula

$$\mathcal{M}_{dynamic} = \llbracket pe(I_t, \mathcal{F}(*, T)) > \eta \cdot pe(I_t, \mathcal{F}(*, T^F)) \rrbracket \odot \llbracket ge(*, T) > \eta \cdot ge(*, T^F) \rrbracket \quad (10)$$

where $*$ represents other inputs required, and η is a scale factor for robust and stable selection, which we set as 1.2 to prefer the ego-motion. Then we can modify the above consistency losses as

$$\tilde{\mathcal{E}}(*, T, T^F) = \mathcal{M}_{dynamic} \odot \mathcal{E}(*, T^F) + (1 - \mathcal{M}_{dynamic}) \odot \mathcal{E}(*, T), \mathcal{E} \in \{pe, ge\} \quad (11)$$

$$\mathcal{L}_{pe}^M = \mathcal{M}_{auto} \odot \min_{t'} \tilde{pe}(*, T, T^F) \quad (12)$$

$$\mathcal{L}_{ge}^M = \mathcal{M}_{auto} \odot \min_{t'} \tilde{ge}(*, T, T^F) \quad (13)$$

3.4 Multi-Phase training

To effectively utilize the proposed ASANet and dynamic-aware learning, we train our networks in three phases, inspired by [23]. For clear expression, we use Θ_D , Θ_{ASA} , Θ_E , Θ_F to represent the DepthNet, the ASANet-encoder, the ego-motion decoder, and the motion field decoder respectively, all of which make up our total framework $\Xi = \{\Theta_D, \Theta_{ASA}, \Theta_E, \Theta_F\}$. And we use these notations as the input parameters to indicate which parts are optimized by a specific objective.

In the first phase, we set aside the motion field decoder of MotionNet, and learn the depth and the camera’s ego-motion that establish the rigid global projection for the warping operation. The overall loss in this phase is a weighted sum of photometric, geometric and disparity smoothness loss.

$$\mathcal{L}_1 = \mathcal{L}_{pe}(\Theta_D, \Theta_{ASA}, \Theta_E) + \lambda_G \mathcal{L}_{ge}(\Theta_D, \Theta_{ASA}, \Theta_E) + \lambda_D \mathcal{L}_{ds}(\Theta_D) \quad (14)$$

After initially training the DepthNet and the static path of MotionNet, we freeze them in the second phase and optimize the motion field decoder. We use the estimated motion field to compute a dynamic transformation for each pixel to handle moving objects. Here, we emphasize again that our proposed ASANet can extract both static and dynamic features, the latter of which is used in this phase.

$$\mathcal{L}_2 = \mathcal{L}_{pe}(\Theta_F) + \lambda_G \mathcal{L}_{ge}(\Theta_F) + \lambda_F \mathcal{L}_{fs}(\Theta_F) \quad (15)$$

In the last phase, we jointly training all components in our system. We merge the two optimization paths in prior phases through the auto-selecting mechanism described in Sec. 3.2, rather than simply glue them together. Therefore, the objective is defined as

$$\mathcal{L}_3 = \mathcal{L}_{pe}^M(\Xi) + \lambda_G \mathcal{L}_{ge}^M(\Xi) + \lambda_F \mathcal{L}_{fs}(\Theta_{ASA}, \Theta_F) + \lambda_D \mathcal{L}_{ds}(\Theta_D) \quad (16)$$

¹You can find more visual results at <https://youtu.be/IMBLIQ5NDUU>.

Method	Sup.	Resolution	Dataset	Error metric ↓				Accuracy metric ↑		
				Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Eigen [2]	D	576 x 172	K	0.203	1.548	6.307	0.282	0.702	0.890	0.890
Kuznietsov [31]	D+S	621 x 187	K	0.113	0.741	4.621	0.189	0.862	0.960	0.986
DORN [5]	D	512 x 385	K	0.072	0.307	2.727	0.120	0.932	0.984	0.995
SfMLearner [10]	M	416 x 128	CS+K	0.198	1.836	6.565	0.275	0.718	0.901	0.960
DF-Net [32]	M	576 x 160	CS+K	0.146	1.182	5.215	0.213	0.818	0.943	0.978
Struct2Depth [14]	M	416 x 128	K	0.141	1.026	5.291	0.215	0.816	0.945	0.979
CC [33]	M	832 x 256	CS+K	0.139	1.032	5.199	0.213	0.827	0.943	0.977
SC-SfMLearner [30]	M	832 x 256	CS+K	0.128	1.047	5.234	0.208	0.846	0.947	0.976
Self-Mono-SF [24]	M	832 x 256	K	0.125	0.978	4.877	0.208	0.851	0.950	0.978
Monodepth2 (R18) [13]	M	640 x 192	K	0.115	0.903	4.863	0.193	0.877	0.959	0.981
TrianFlow [22]	M	832 x 256	K	0.113	0.704	4.581	0.184	0.871	0.961	0.984
Monodepth2 (R50) [13]	M	640 x 192	K	0.111	0.825	4.644	0.187	0.883	0.962	0.982
PackNet-SfM [34]	M	640 x 192	K	0.111	0.785	4.601	0.189	0.878	0.960	0.982
Ours (R18)	M	640 x 192	K	0.112	0.866	4.693	0.189	0.881	0.961	0.981
Ours (R50)	M	640 x 192	K	0.108	0.820	4.595	0.186	0.886	0.963	0.982

Table 1: **Quantitative results on the Eigen split of KITTI dataset** for distances up to 80m. For training dataset, K means that the method is trained on the KITTI dataset [18], while CS+K refers to the models with pre-training on Cityscapes [35]. D denotes the depth supervision, and S denotes stereo input pairs. M refers to the methods in the self-supervised learning manner with monocular video clips. (R18) and (R50) refer to the models with ResNet18 or ResNet50 as the backbone. Best results in supervised and self-supervised methods are in bold for highlight.

4 Experiments

To validate the contributions of our method, we conduct extensive experiments and describe the results in this section. We train and test our models on the KITTI dataset [18] for depth estimation and on the KITTI odometry dataset [19] for visual odometry, for the fair comparison with existing self-supervised methods.

4.1 Datasets

KITTI dataset The KITTI dataset [18] is a standard benchmark for evaluating depth prediction, with LiDAR readings as the ground-truths for evaluation. We use the data split of Eigen et al. [2] and adopt the pre-processing of Zhou et al. [10] to remove stationary frames. This results in 39810 snippets for training, 4424 for validation, and 697 for testing.

KITTI Odometry dataset KITTI Odometry dataset [19] contains 22 stereo sequences with 11 sequences having accessible ground-truth camera poses. Consistently with the literature [10, 13, 22, 25], we train our models on sequences 00-08 and evaluate on sequences 09 and 10 with the length of 1591 and 1201 respectively. Then we have 36671 training triplets and 4075 for validation.

4.2 Implementation details

In our framework, the DepthNet is a U-shaped architecture as in [13]. The MotionNet is as described in Sec. 3.1. The two networks both use ResNet [36] as the backbone which is pretrained on ImageNet [27].

Our models are implemented in PyTorch [37]. For all experiments, we train our models with 3-frame sequences $\{I_{t-1}, I_t, I_{t+1}\}$ in one step, and resize the input images to 640×192 during both training and testing. We use the Adam optimizer [38] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The loss weightings are $[\lambda_G, \lambda_F, \lambda_D] = [0.1, 0.001, 0.001]$. To evaluate the performance of our proposed framework, we train our models with the multi-phase schedule (Sec. 3.4), where each phase is for 10 epochs. Since the optimized components in the first two phases have no overlap, we train each component in our system for 20 epochs in total with a batch size of 4. We set the initial learning rate as 10^{-4} , which is decayed by 10 for refinement after 15 total training epochs.

4.3 Evaluation of depth estimation

For depth estimation, we train our models in the three-stage schedule on preprocessed KITTI dataset [18], as described in Sec. 3.4. Following Godard et al. [13], we capture the depth up to 80m and scale the depth to align the per-image median with the median of ground-truth during evalua-

Type	Method	Latency	Seq. 09		Seq. 10	
			$t_{err}(\%)$	$r_{err}(o/100m)$	$t_{err}(\%)$	$r_{err}(o/100m)$
Geometric VO	ORB-SLAM2 (w/o LC) [39]	-	9.31	0.26	2.66	0.39
	ORB-SLAM2 (w/ LC) [39]	-	2.84	0.25	2.67	0.38
	TrianFlow [22] [‡]	120.25 ms	6.93	0.44	4.66	0.62
End-to-end VO	SfMLearner [10]	1.68 ms	11.32	4.07	15.25	4.06
	Depth-VO-Feat [40]	2.11 ms	9.07	3.80	9.60	3.41
	SC-SfMLearner [30] [†]	9.38 ms	7.64	2.19	10.74	4.58
	Ours (R18)	20.68 ms	5.33	2.13	8.79	3.46

Table 2: **Quantitative results on KITTI Odometry Sequences 09 and 10.** The latency for one inference, average translation (t_{err}), and rotation errors (r_{err}) are reported. [†] indicates that we use their updated results in [25]. [‡] indicates when testing its running time, we only report the time it uses a PWC-Net [41] to infer the forward and backward optical flow. LC=Loop Closure.

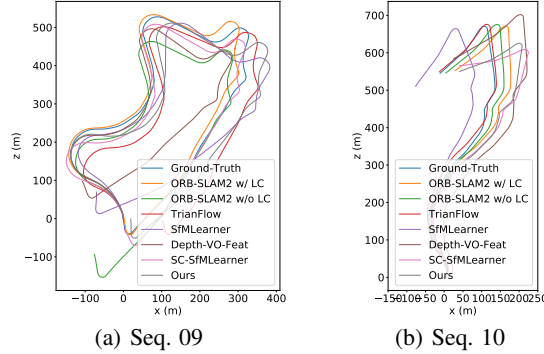


Figure 3: **Qualitative trajectory results on KITTI Odometry Sequences 09 and 10.**

tion. We compare our method with the existing methods in Tab. 1, which shows that our method can achieve state-of-the-art performance. Despite a performance gain for depth estimation, we use the standard U-shaped DepthNet without any additional computation cost, which will be illustrated in Sec. 4.5. That’s to say, if we incorporate an improved DepthNet, such as PackNet proposed in [34], we can further boost the performance of our method. In other words, our method can be integrated with other self-supervised methods by replacing the traditional PoseNet for ego-motion estimation by our MotionNet and adopting the dynamic-aware auto-selecting mechanism during training.

To verify the effect of our proposed auto-selecting mechanism, we visualize the qualitative samples in Fig. 2, consisting of sample images, estimated depth, masks (\mathcal{M}_{auto} and $\mathcal{M}_{dynamic}$) from auto-masking and auto-selecting approaches, and photometric error maps before and after fusing the rigid global projection with the dynamic pixel-wise projection. From the third row in Fig. 2, we can see that the auto-masking \mathcal{M}_{auto} can reduce the influence of the sky and other areas where the texture change is not obvious, while our dynamic-aware mask $\mathcal{M}_{dynamic}$ can effectively identify the moving objects in the scene. The change of the photometric error maps in the last line can explain that through our auto-selecting mechanism, the redundant errors caused by moving objects can be reduced, thereby avoiding the negative effects of these outliers. Besides, for areas with more detailed textures, such as leaf and bush in Samples 1 and 2, the auto-selecting mechanism can also provide better supervision.

4.4 Evaluation of visual odometry

We use the same experimental settings as the depth estimation and train our networks on the KITTI odometry sequences 00-08 according to only the first phase in Sec. 3.4, to compare our methods with some recent self-supervised VO methods on sequences 09 and 10. We scale and use a 7-DoF transformation to align our estimated poses with the ground-truth as Zhan et al. [25] did. As a robot localization module that needs to be run in real-time, the model’s run-time latency has a great influence on the practical application. So we compare our methods with the existing methods on both performance and the run-time latency in Tab. 2. To test the latency of the model, we forward each model a thousand times on an Nvidia GeForce RTX 2080Ti GPU and take the average running time as the results. It is worth mentioning that in the inference process of VO, we only need to run ASANet-encoder and ego-motion decoder of our MotionNet. Especially, TrianFlow [22] is a

Method	#Params	Abs Rel	Sq Rel	$\delta_{1.25}$
PackNet-SfM (D=4) [34]	78.49 M	0.113	0.818	0.875
PackNet-SfM (D=8) [34]	128.29 M	0.111	0.785	0.878
Monodepth2 (R18) [13]	14.84 M	0.115	0.903	0.877
+ \mathcal{L}_{ge} (Eq. (6))	14.84 M	0.114	0.874	0.877
+ \mathcal{L}_{ge} + ASANet (Sec. 3.1) (ours)	14.84 M	0.113	0.870	0.878
+ \mathcal{L}_{ge} + ASANet + Auto-selecting (Secs. 3.3 and 3.4) (ours)	14.84 M	0.112	0.867	0.882
Monodepth2 (R50) [13]	34.57 M	0.111	0.825	0.883
+ \mathcal{L}_{ge} (Eq. (6))	34.57 M	0.109	0.823	0.885
+ \mathcal{L}_{ge} + ASANet (Sec. 3.1) (ours)	34.57 M	0.109	0.824	0.886
+ \mathcal{L}_{ge} + ASANet + Auto-selecting (Secs. 3.3 and 3.4) (ours)	34.57 M	0.108	0.820	0.886

Table 3: **Ablation study** on standard KITTI benchmark [18] which compares the proposed methods with our baseline, Monodepth2 [13], and a state-of-the-art method, PackNet-SfM [34]. (R18) and (R50) indicates the backbone of the DepthNet.

two-stage method that first uses PWC-Net [41] to estimate forward and backward optical flow, then utilizes the geometric VO methods. In our experiments, we only report the running time for its two inferences for forward and backward optical flow estimation. Compared with ORB-SLAM2 [39], TriangFlow uses FlowNet to replace the process of calculating geometric feature descriptors and matching points, which is still a geometry-based VO, so we group them together. It can be seen from the comparison in Tab. 2 that our method can achieve the best in one-stage end-to-end VO methods, and maintain a real-time operating speed of 50 fps. The visualized trajectories on KITTI Odometry sequences 09 and 10 are shown in Fig. 3.

4.5 Ablation study

The performance of our novel self-supervised joint learning system has been shown above. In this section we want to verify the effects of the several improvements we have proposed, including the ASANet (Sec. 3.1) and the dynamic-aware learning (Secs. 3.3 and 3.4). We take the Monodepth2 [13] as a strong baseline, and the state-of-the-art PackNet-SfM [34] as a strong reference. In the experimental setting without the auto-selecting mechanism, we train our models according to the loss function shown in Eq. (14), and extend its training time to 20 epochs. For the complete dynamic-aware method that incorporates both ASANet and the auto-selecting mechanism, we use the multi-stage schedule proposed in Sec. 3.4 for training, and each stage lasts for 10 epochs.

The experimental results are shown in Tab. 3, which demonstrates that the system’s performance can be improved by adding our proposed ASANet and multi-phase dynamic-aware learning. Compared with PackNet-SfM [34], our method can boost the system’s performance while the model parameters are less than half of it. This makes our model more advantageous when deployed on the robot. Besides, we notice that when we use ResNet50 as the DepthNet’s backbone, the gain effect of ASANet without auto-selecting is reduced in the depth estimation task, but the auto-selecting mechanism for dynamic objects can still improve the performance of the model.

5 Conclusion

In this paper, we mainly track moving objects issue, which is a stumbling block in self-supervised video learning. Unlike many previous methods [14, 15, 23] that use additional information such as segmentation and stereo cues, we propose an Attentional Separation-and-Aggregation Network (ASANet) that can automatically distinguish dynamic and static features in the scene, and design an auto-selecting mechanism to achieve dynamic-aware learning. We conduct extensive experiments to verify the performance of our method. The results show that the auto-selecting mechanism we propose can automatically detect moving objects and select appropriate transformations for reconstruction, thus avoiding the interference of these exceptions to the joint depth-pose learning. Moreover, our method can achieve competitive performance on depth estimation and visual odometry while maintaining a low computational cost compared with the existing state-of-the-art methods. In future work, we can try further to optimize the system from the perspective of depth estimation, such as incorporating the 3D-Packing proposed in PackNet-SfM [34], or by utilizing the RNN structure often used in supervised visual odometry methods [20, 21] to leverage the historical information in the video.

Acknowledgments

This work is supported by National Key R&D Program of China (2018YFB0105000). This work is also supported by Meituan-Dianping Group, National Natural Science Foundation of China (No. U19B2019, 61832007, 61621091), Tsinghua EE Xilinx AI Research Fund, Beijing National Research Center for Information Science and Technology (BNRist) and Beijing Innovation Center for Future Chips.

References

- [1] M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, June 1981. ISSN 0001-0782.
- [2] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in Neural Information Processing Systems 27*, pages 2366–2374. 2014.
- [3] D. Xu, E. Ricci, W. Ouyang, X. Wang, and N. Sebe. Multi-scale continuous crfs as sequential deep networks for monocular depth estimation. In *CVPR*, 2017.
- [4] D. Xu, W. Wang, H. Tang, H. Liu, N. Sebe, and E. Ricci. Structured attention guided convolutional neural fields for monocular depth estimation. In *CVPR*, 2018.
- [5] H. Fu, M. Gong, C. Wang, K. Batmanghelich, and D. Tao. Deep Ordinal Regression Network for Monocular Depth Estimation. In *CVPR*, 2018.
- [6] A. Kendall, M. Grimes, and R. Cipolla. PoseNet: A convolutional network for real-time 6-dof camera relocalization. In *ICCV*, pages 2938–2946, 2015. ISBN 978-1-4673-8391-2.
- [7] S. Wang, R. Clark, H. Wen, and A. Trigoni. Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks. *ICRA*, pages 2043–2050, 2017.
- [8] J. Jiao, J. Jiao, Y. Mo, W. Liu, and Z. Deng. Magicvo: An end-to-end hybrid cnn and bi-lstm method for monocular visual odometry. *IEEE Access*, 7:94118–94127, 2019.
- [9] M. R. U. Saputra, P. Gusmao, S. Wang, A. Markham, and N. Trigoni. Learning monocular visual odometry through geometry-aware curriculum learning. In *ICRA*, 03 2019.
- [10] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR 2017*, pages 6612–6619.
- [11] R. Mahjourian, M. Wicke, and A. Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In *CVPR*, 2018.
- [12] J.-W. Bian, Z. Li, N. Wang, H. Zhan, C. Shen, M.-M. Cheng, and I. Reid. Unsupervised scale-consistent depth and ego-motion learning from monocular video. In *NeurIPS*, 2019.
- [13] C. Godard, O. Mac Aodha, M. Firman, and G. Brostow. Digging into self-supervised monocular depth estimation. In *ICCV 2019*.
- [14] V. Casser, S. Pirk, R. Mahjourian, and A. Angelova. Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:8001–8008, 07 2019.
- [15] A. Gordon, H. Li, R. Jonschkowski, and A. Angelova. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *ICCV*, October 2019.
- [16] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017.
- [17] S. Li, X. Wang, Y. Cao, F. Xue, Z. Yan, and H. Zha. Self-supervised deep visual odometry with online adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

- [18] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013.
- [19] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [20] J. Jiao, Y. Cao, Y. Song, and R. Lau. Look deeper into depth: Monocular depth estimation with semantic booster and attention-driven loss. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [21] 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 *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 8575–8583, 2019.
- [22] W. Zhao, S. Liu, Y. Shu, and Y.-J. Liu. Towards better generalization: Joint depth-pose learning without posenet. In *CVPR*, 2020.
- [23] S. Lee, S. Im, S. Lin, and I. S. Kweon. Learning residual flow as dynamic motion from stereo videos. *arXiv preprint arXiv:1909.06999*, 2019.
- [24] J. Hur and S. Roth. Self-supervised monocular scene flow estimation. In *CVPR*, 2020.
- [25] H. Zhan, C. S. Weerasekera, J. Bian, and I. Reid. Visual odometry revisited: What should be learnt? In *ICRA 2020*.
- [26] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 770–778, 2016.
- [27] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009.
- [28] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 7132–7141, 2018.
- [29] Zhou Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004.
- [30] J.-W. Bian, Z. Li, N. Wang, H. Zhan, C. Shen, M.-M. Cheng, and I. Reid. Unsupervised scale-consistent depth and ego-motion learning from monocular video. In *NeurIPS 2019*.
- [31] Y. Kuznetsov, J. Stuckler, and B. Leibe. Semi-supervised deep learning for monocular depth map prediction. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [32] Y. Zou, Z. Luo, and J.-B. Huang. Df-net: Unsupervised joint learning of depth and flow using cross-task consistency. In *European Conference on Computer Vision*, 2018.
- [33] A. Ranjan, V. Jampani, L. Balles, K. Kim, D. Sun, J. Wulff, and M. J. Black. Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In *CVPR*, pages 12240–12249, 2019.
- [34] V. Guizilini, R. Ambrus, S. Pillai, A. Raventos, and A. Gaidon. 3d packing for self-supervised monocular depth estimation. In *CVPR*, 2020.
- [35] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- [36] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2015.
- [37] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, and etc. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8026–8037. 2019.

- [38] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [39] R. Mur-Artal and J. D. Tardós. ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras. *IEEE Transactions on Robotics*, 33(5):1255–1262, 2017.
- [40] H. Zhan, R. Garg, C. Saroj Weerasekera, K. Li, H. Agarwal, and I. Reid. Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [41] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *CVPR*, 2018.