

DytanVO: Joint Refinement of Visual Odometry and Motion Segmentation in Dynamic Environments

Shihao Shen, Yilin Cai, Wenshan Wang, Sebastian Scherer

Abstract—Learning-based visual odometry (VO) algorithms achieve remarkable performance on common static scenes, benefiting from high-capacity models and massive annotated data, but tend to fail in dynamic, populated environments. Semantic segmentation is largely used to discard dynamic associations before estimating camera motions but at the cost of discarding static features and is hard to scale up to unseen categories. In this paper, we leverage the mutual dependence between camera ego-motion and motion segmentation and show that both can be jointly refined in a single learning-based framework. In particular, we present DytanVO, the first supervised learning-based VO method that deals with dynamic environments. It takes two consecutive monocular frames in real-time and predicts camera ego-motion in an iterative fashion. Our method achieves an average improvement of 27.7% in ATE over state-of-the-art VO solutions in real-world dynamic environments, and even performs competitively among dynamic visual SLAM systems which optimize the trajectory on the backend. Experiments on plentiful unseen environments also demonstrate our method’s generalizability.

I. INTRODUCTION

Visual odometry (VO), one of the most essential components for pose estimation in the visual Simultaneous Localization and Mapping (SLAM) system, has attracted significant interest in robotic applications over past few years [1]. A lot of research works have been conducted to develop an accurate and robust monocular VO system using both geometry-based methods [2], [3]. However, they require significant engineering effort for each module to be carefully designed and finetuned [4], which makes it difficult to be readily deployed in the open world with complex environmental dynamics, changes of illumination or inevitable sensor noises.

On the other hand, recent learning-based methods [4]–[7] are able to outperform geometry-based methods in more challenging environments such as large motion, fog or rain effects and lack of features. However, they will easily fail in dynamic environments if they do not take into consideration independently moving objects that cause unpredictable changes in illumination or occlusions. To this end, recent works utilize abundant unlabeled data and adopt either self-supervised learning [8], [9] or unsupervised learning [10], [11] to handle dynamic scenes. Although they achieve outstanding performance on particular tasks, such as autonomous driving, they produce worse results if applied to very different data distributions, such as micro air vehicles

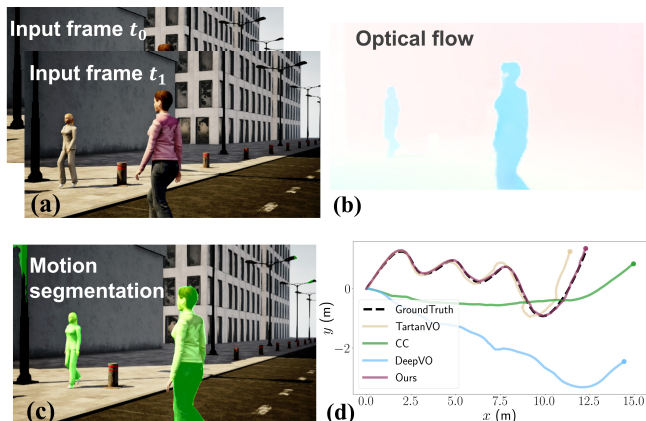


Fig. 1: A overview of the DytanVO. (a) Input frame at time t_0 and t_1 . (b) Optical flow output from the matching network. (c) Motion segmentation output after iterations. (d) Trajectory estimation on sequence RoadCrossing VI from the AirDOS-Shibuya Dataset, which is a highly dynamic environment cluttered with humans. Ours is the only learning-based VO that keeps track.

(MAV) that operate with aggressive and frequent rotations cars do not have. Learning without supervision is hindered from generalizing due to biased data with simple motion patterns. Therefore, we approach the dynamic VO problem as supervised learning so that the model can map inputs to complex ego-motion ground truth and be more generalizable.

To identify dynamic objects, object detection or semantic segmentation techniques are largely relied on to mask all movable objects, such as pedestrians and vehicles [12]–[15]. Their associated features are discarded before applying geometry-based methods. However, there are two issues of utilizing semantic information in dynamic VO. First, class-specific detectors for semantic segmentation heavily depend on appearance cues but not every object that can move is present in the training categories, leading to false negatives. Second, even if all moving objects in a scene within the categories, algorithms could not distinguish between “actually moving” versus “static but being able to move”. In dynamic VO where static features are crucial to robust ego-motion estimation, one should segment objects based on pure motion (motion segmentation) rather than heuristic appearance cues.

Motion segmentation utilizes relative motion between consecutive frames to remove the effect of camera movement from the 2D motion fields and calculates residual optical flow to account for moving regions. But paradoxically, ego-motion cannot be correctly estimated in dynamic scenes without a robust segmentation. There exists such a mutual dependence

Code is available at <https://github.com/Geniush/DytanVO>

S. Shen, Y. Cai, W. Wang, S. Scherer are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA. {shihaosh, yilinc, wenshanw, basti}@andrew.cmu.edu

between motion segmentation and ego-motion estimation that has never been explored in supervised learning methods. Therefore, motivated by jointly refining the VO and motion segmentation, we propose our learning-based dynamic VO (DyTanVO). To our best knowledge, our work is the first supervised learning-based VO for dynamic environments. The main contributions of this paper are threefold:

- A novel learning-based VO is introduced to leverage the interdependence among camera ego-motion, optical flow and motion segmentation.
- We introduce an iterative framework where both ego-motion estimation and motion segmentation can converge quickly within time constraints for real-time applications.
- Among learning-based VO solutions, our method achieves state-of-the-art performance in real-world dynamic scenes without finetuning. Furthermore, our method performs even comparably with visual SLAM solutions that optimize trajectories on the backend.

II. RELATED WORK

Learning-based VO solutions aim to avoid hard-coded modules that require significant engineering efforts for design and finetuning in classic pipelines [1], [16]. For example, Valada [17] applies auxiliary learning to leverage relative pose information to constrain the search space and produce consistent motion estimation. Another class of learning-based methods rely on dense optical flow to estimate pose as it provides more robust and redundant modalities for feature association in VO [5], [18], [19]. However, their frameworks are built on the assumption of photometric consistency which only holds in a static environment without independently moving objects. They easily fail when dynamic objects unpredictably cause occlusions or illuminations change.

Semantic information is largely used by earlier works in VO or visual SLAM to handle dynamic objects in the scene, which is obtained by either a feature-based method or a learning-based method. Feature-based methods utilize hand-designed features to recognize semantic entities [20]. An exemplary system proposed by [21] computes SIFT descriptors from monocular image sequences in order to recognize semantic objects. On the other hand, data-driven CNN-based semantic methods have been widely used to improve the performance, such as DS-SLAM [22] and SemanticFusion [23]. A few works on semantic VO/SLAM have fused the semantic information from recognition modules to enhance motion estimation and vice versa [24], [25]. However, all these methods are prone to limited semantic categories, which leads to false negatives when scaling to unusual real-world applications such as offroad driving or MAV, and requires continuous efforts in ground-truth labeling.

Instead of utilizing appearance cues for segmentation, efforts are made to segment based on geometry cues. Flow-Fusion [26] iteratively refines its ego-motion estimation by computing residual optical flow. GeoNet [10] divides its system into two sub-tasks by separately predicting static scene structure and dynamic motions. However, both depend

on geometric constraints arising from epipolar geometry and rigid transformations, which are vulnerable to motion ambiguities such as objects moving in the colinear direction relative to the camera being indistinguishable from the background given only ego-motion and optical flow. On the other hand, MaskVO [8] and SimVODIS++ [9] approach the problem by learning to mask dynamic feature points in a self-supervised manner. CC [11] couples motion segmentation, flow, depth and camera motion models which are jointly solved in an unsupervised way. Nevertheless, these self-supervised or unsupervised methods are trained on self-driving vehicle data dominated by pure translational motions with little rotation, which makes them difficult to generalize to completely different data distributions such as handheld cameras or drones. Our work introduces a framework that jointly refines camera ego-motion and motion segmentation in an iterative way that is robust against motion ambiguities as well as generalizes to the open world.

III. METHODOLOGY

A. Datasets

Built on TartanVO [5], our method remains its generalization capability while handling dynamic environments in multiple types of scenes, such as car, MAV, indoor and outdoor. Besides taking camera intrinsics as an extra layer into the network to adapt to various camera settings as explored in [5], we train our model on large amounts of synthetic data with broad diversity, which is shown capable of facilitating easy adaptation to the real world [27]–[29].

Our model is trained on both TartanAir [27] and SceneFlow [30]. The former contains more than 400,000 data frames with ground truth of optical flow and camera pose in static environments only. The latter provides 39,000 frames in highly dynamic environments with each trajectory having backward/forward passes, different objects and motion characteristics. Although SceneFlow does not provide ground truth of motion segmentations, we are able to recover it by taking use of its ground truth of disparity, optical flow and disparity change maps.

B. Architecture

Our network architecture is illustrated in Fig. 2, which is based on TartanVO. Our method takes in two consecutive undistorted images I_t, I_{t+1} and outputs the relative camera motion $\delta_t^{t+1} = (\mathbf{R}|\mathbf{T})$, where $\mathbf{T} \in \mathbf{R}^3$ is the 3D translation and $\mathbf{R} \in SO(3)$ is the 3D rotation. Our framework consists of three sub-modules, a matching network, a motion segmentation network, and a pose network. We estimate dense optical flow F_t^{t+1} with a matching network, $M_\theta(I_t, I_{t+1})$, from two consecutive images. The network is built based on PWC-Net [31]. The motion segmentation network U_γ , based on a lightweight U-Net [32], takes in the relative camera motion output, $\mathbf{R}|\mathbf{T}$, optical flow from M_θ , and the original input frames. It outputs a probability map, z_t^{t+1} , of every pixel belonging to a dynamic object or not, which is thresholded and turned into a binary segmentation mask, S_t^{t+1} . The optical flow is then stacked with the mask and

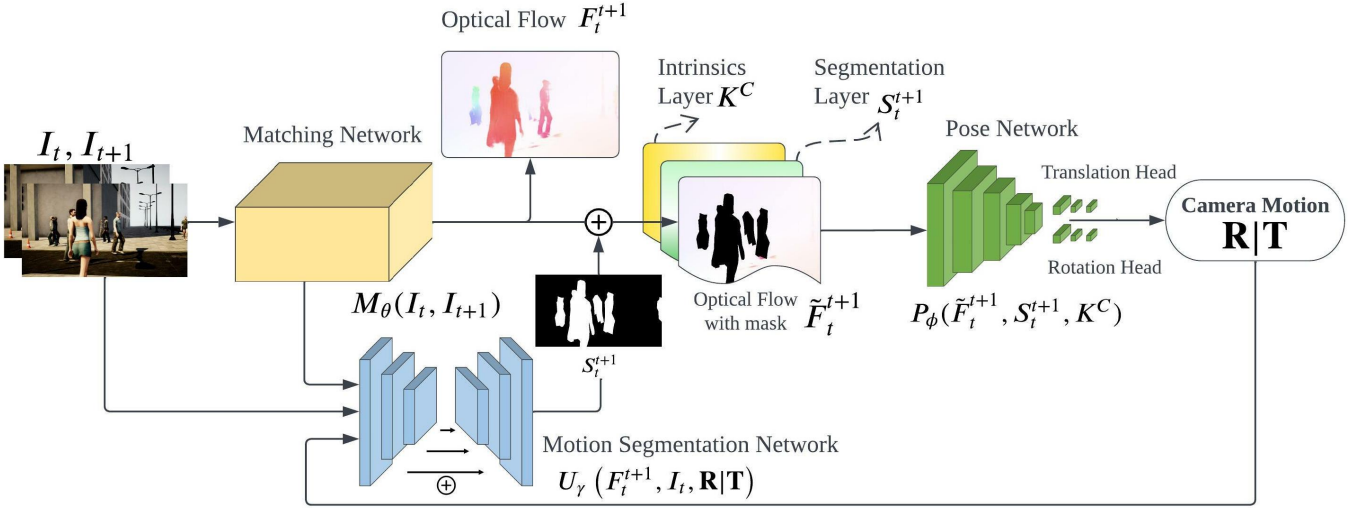


Fig. 2: Overview of our three-stage network architecture. It consists of a matching network which estimates optical flow from two consecutive images, a pose network that estimates pose based on optical flow without dynamic movements, and a motion segmentation network that outputs a probability mask of the dynamicness. The matching network is forwarded only once while the pose network and the segmentation network are iterated to jointly refine pose estimate and motion segmentation. In the first iteration, we randomly initialize the segmentation mask. In each iteration, optical flow is set to zero inside masked regions.

the intrinsic layer K^C , followed by setting all optical flow inside the masked regions to zeros, i.e., \tilde{F}_t^{t+1} . The last is a pose network P_ϕ , with ResNet50 [33] as the backbone, which takes in the previous stack, and outputs camera motion.

C. Motion segmentation

Earlier dynamic VO methods that use motion segmentation rely on purely geometric constraints arising from epipolar geometry and rigid transformations [12], [26] so that they can threshold residual optical flow which is designed to account for moving regions. However, they are prone to catastrophic failures under two cases: (1) points in 3D moving along epipolar lines cannot be identified from the background given only monocular cues; (2) pure geometry methods leave no tolerance to noisy optical flow and less accurate camera motion estimations, which in our framework is very likely to happen in the first few iterations. Therefore, following [34], to deal with the ambiguities above, we explicitly model cost maps as inputs into the segmentation network after upgrading the 2D optical flow to 3D through optical expansion [35], which estimates the relative depth based on the scale change of overlapping image patches. The cost maps are tailored to coplanar and colinear motion ambiguities that cause segmentation failures in geometry-based motion segmentation. More details can be found in [34].

D. Iteratively refine camera motion

We provide an overview of our iterative framework in Algorithm 1. During inference, the matching network is forwarded only once while the pose network and the segmentation network are iterated to jointly refine ego-motion estimation and motion segmentation. In the first iteration, the segmentation mask is initialized randomly using [36]. The criterion to stop iteration is straightforward, which is the rotational and translational differences of $\mathbf{R}|\mathbf{T}$ between two

iterations being smaller than prefixed thresholds ϵ . Instead of having a fixed constant to threshold probability maps into segmentation masks, we predetermine a decaying parameter that empirically reduces the input threshold over time, in order to discourage inaccurate masks in earlier iterations while embracing refined masks in later ones.

Algorithm 1 Inference with Iterations

Given two consecutive frames I_t, I_{t+1} and intrinsic K
Initialize iteration number: $i \leftarrow 1$
Initialize difference in output camera motions: $\delta_{R|T} \leftarrow \infty$
 ${}^i F_t^{t+1} \leftarrow \text{OpticalFlow}(I_t, I_{t+1})$
while $\delta_{R|T} \geq \text{stopping criterion}, \epsilon$ **do**
 if i is 0 **then**
 ${}^i S_t^{t+1} \leftarrow \text{getCowmask}(I_t)$
 else
 ${}^i z_t^{t+1} \leftarrow \text{MotionSegmentation}({}^i F_t^{t+1}, I_t, {}^i \mathbf{R}|\mathbf{T})$
 ${}^i S_t^{t+1} \leftarrow \text{mask } {}^i z_t^{t+1} \geq z_{\text{threshold}}$
 ${}^i \tilde{F}_t^{t+1} \leftarrow \text{set } {}^i F_t^{t+1} = 0 \text{ for } {}^i S_t^{t+1} == 1$
 ${}^i \mathbf{R}|\mathbf{T} \leftarrow \text{PoseNetwork}({}^i \tilde{F}_t^{t+1}, {}^i S_t^{t+1}, K)$
 $\delta_{R|T} \leftarrow {}^i \mathbf{R}|\mathbf{T} - {}^{i-1} \mathbf{R}|\mathbf{T}$
 $i \leftarrow i + 1$

Intuitively, during early iterations, the estimated motion is less accurate, which leads to false positives in the segmentation output (assigning high probabilities to static areas). However, due to the fact that optical flow map still provides enough correspondences regardless of cutting out non-dynamic regions from it, P_ϕ is able to robustly leverage the segmentation mask S_t^{t+1} concatenated with \tilde{F}_t^{t+1} , and outputs reasonable camera motion. In later iterations, U_γ is expected to output increasingly precise probability maps such that static regions in the optical flow map are no longer

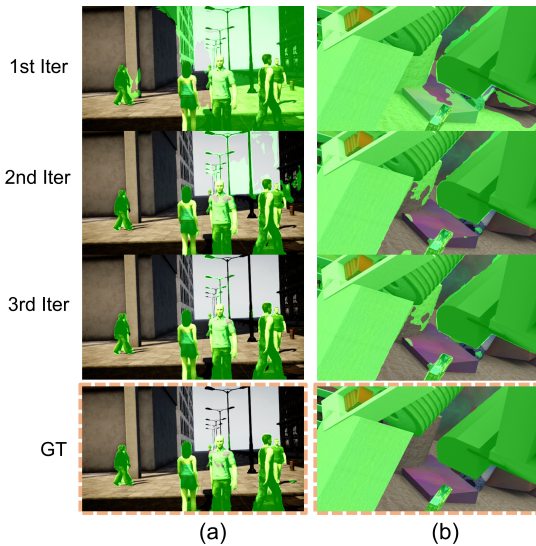


Fig. 3: Motion segmentation output at each iteration when testing on unseen data. (a) Running inference on the hardest sequence in AirDOS-Shibuya with multiple people moving in different directions with our segmentation network. (b) Inference on the sequence from FlyingThings3D where dynamic objects take up more than 60% area. Ground truth (GT) mask on Shibuya is generated by the segmentation network with GT ego-motion as input.

“wasted” and hence P_ϕ can be improved accordingly.

In practice, we find that 3 iterations are more than enough to get both camera motion and segmentation refined. To clear up any ambiguity, a 1-iteration pass is composed of one M_θ forward pass and one P_ϕ forward pass with random mask, while a 3-iteration pass consists of one M_θ forward pass, two U_γ forward passes and three P_ϕ forward passes. In Fig. 3 we illustrate how segmentation masks evolve after three iterations on unseen data. The mask at the first iteration contains a significant amount of false positives but quickly converges beyond the second iteration. This verifies our assumption that the pose network is robust against false positives in segmentation results.

E. Supervision

We train our pose network to be robust against large areas of false positives. On training data without any dynamic object, we adopt the cow-mask [36] to create sufficiently random yet locally connected segmentation patterns as a motion segmentation could occur in any size, any shape and at any position in an image while exhibiting locally explainable structures corresponding to the types of moving objects. In addition, we apply curriculum learning to the pose network where we gradually increase the maximum percentage of dynamic areas in SceneFlow from 15%, 20%, 30%, 50% to 100%. Since TartanAir only contains static scenes, we adjust the size of the cow-masks accordingly.

We supervise our network on the camera motion loss L_P . Under the monocular setting, we only recover an up-to-scale camera motion. We follow [5] and normalize the translation vector before calculating the distance to ground truth. Given

ground truth motion $R|T$,

$$L_P = \left\| \frac{\hat{\mathbf{T}}}{\max(\|\hat{\mathbf{T}}\|, \epsilon)} - \frac{\mathbf{T}}{\max(\|\mathbf{T}\|, \epsilon)} \right\| + \|\hat{\mathbf{R}} - \mathbf{R}\| \quad (1)$$

where $\epsilon=1e-6$ to prevent numerical instability and $\hat{\cdot}$ denotes estimated quantities.

Our framework can also be trained in an end-to-end fashion, in which case the objective becomes an aggregated loss of the optical flow loss L_M , the camera motion loss L_P and the motion segmentation loss L_U , where L_M is the L1 norm between the predicted flow and the ground truth flow whereas L_U is the binary cross entropy loss between predicted probability and the segmentation label.

$$L = \lambda_1 L_M + \lambda_2 L_U + L_P \quad (2)$$

From preliminary empirical comparison, end-to-end training gives similar performance to training the pose network only, because we use λ_1 and λ_2 to regularize the objective such that the training is biased toward mainly improving the odometry rather than optimizing the other two tasks. This is ideal since the pose network is very tolerant of false positives in segmentation results (shown in III-D). In the following section, we show our results of supervising only on Eq. 1 by fixing the motion segmentation network.

IV. EXPERIMENTAL RESULTS

A. Implementation details

1) *Network*: We initialize the matching network M_θ with the pre-trained model from TartanVO [5], and fix the motion segmentation network U_γ with the pre-trained weights from Yang et al. [34]. The pose network P_ϕ uses ResNet50 [33] as the backbone, removes the batch normalization layers, and adds two output heads for rotation R and translation T . M_θ outputs optical flow at size of $H/4 \times W/4$. P_ϕ takes in a 5-channel input, i.e., $\tilde{F}_t^{t+1} \in \mathbb{R}^{2 \times H/4 \times W/4}$, $S_t^{t+1} \in \mathbb{R}^{H/4 \times W/4}$ and $K^C \in \mathbb{R}^{2 \times H/4 \times W/4}$. The concatenation of \tilde{F}_t^{t+1} and K^C augments the optical flow input with 2D positional information while concatenating \tilde{F}_t^{t+1} with S_t^{t+1} encourages the network to learn dynamic representations.

2) *Training*: Our method is implemented in PyTorch [43] and trained on 2 NVIDIA A100 Tensor Core GPUs. We train the network in two stages on TartanAir, which includes only static scenes, and SceneFlow [30]. In the first stage, we train P_ϕ independently using ground truth optical flow, camera motion, and motion segmentation mask in a curriculum-learning fashion. We generate random cow-masks [36] on TartanAir as motion segmentation input. Each curriculum is initialized with weights from the previous curriculum and takes 100,000 iterations with a batch size of 256. In the second stage, P_ϕ and M_θ are jointly optimized for another 100,000 iterations with a batch size of 64. During curriculum learning, the learning rate starts at $2e-4$, while the second stage uses a learning rate of $2e-5$. Both stages apply a decay rate of 0.2 to the learning rate every 50,000 iterations. Random cropping and resizing (RCR) [5] as well as frame skipping are applied to both datasets.

TABLE I: ATE (m) results on dynamic sequences from AirDOS-Shibuya. Our method gives outstanding performance among VO methods. DeepVO, TrianFlow and CC are trained on KITTI only and unable to generalize to complex motion patterns. All SLAM methods use bundle adjustment (BA) on multiple frames to optimize the trajectory and hence we only numerically compare ours with pure VO methods. The best and the second best VO performances are highlighted as bold and underlined. We use “-” to denote SLAM methods that fail to initialize.

		StandingHuman		RoadCrossing (Easy)			RoadCrossing (Hard)	
		I	II	III	IV	V	VI	VII
SLAM method	DROID-SLAM [37]	0.0051	0.0073	0.0103	0.0120	0.2778	0.0253	0.5788
	AirDOS w/ mask [38]	0.0606	0.0193	0.0951	0.0331	0.0206	0.2230	0.5625
	ORB-SLAM w/ mask [39]	0.0788	0.0060	0.0657	0.0196	0.0148	1.0984	0.8476
	VDO-SLAM [40]	0.0994	0.6129	0.3813	0.3879	0.2175	0.2400	0.6628
	DynaSLAM [41]	-	0.8836	0.3907	0.4196	0.4925	0.6446	0.6539
VO method	DeepVO [4]	0.3956	0.6351	0.7788	0.3436	0.5434	0.7223	0.9633
	TrianFlow [42]	0.9743	1.3835	1.3348	1.6172	1.4769	1.7154	1.9075
	CC [11]	0.4527	0.7714	0.5406	0.6345	0.5411	0.8558	1.0896
	TartanVO [5]	<u>0.0600</u>	<u>0.1605</u>	<u>0.2762</u>	<u>0.1814</u>	<u>0.2174</u>	<u>0.3228</u>	<u>0.5009</u>
	Ours	0.0327	0.1017	0.0608	0.0516	0.0755	0.0365	0.0660

3) *Runtime*: Although our method iterates multiple times to refine both segmentation and camera motion, we find in practice that 3 iterations are more than enough due to the robustness of P_ϕ as shown in Fig. 3. On an NVIDIA RTX 2080 GPU, inference takes 40ms with 1 iteration, 100ms with 2 iterations and 160ms with 3 iterations.

4) *Evaluation*: We use the Absolute Trajectory Error (ATE) to evaluate our algorithm against other state-of-the-art methods including both VO and Visual SLAM. We evaluate our method on AirDOS-Shibuya dataset [38] and KITTI Odometry dataset [44]. Additionally, in the supplemental material, we test our method on data collected in a cluttered intersection to demonstrate our method can scale to real-world dynamic scenes competitively.

B. Performance on AirDOS-Shibuya Dataset

We first provide an ablation study of the number of iterations (iter) in Tab. III using three sequences from AirDOS-Shibuya [38]. The quantitative results are consistent with Fig. 3 where the pose network quickly converges after the first iteration. We also compare the 3-iteration finetuned model after jointly optimizing P_ϕ and M_θ (second stage), which shows less improvement because the optical flow estimation on AirDOS-Shibuya already has high quality.

TABLE III: Experiments on number of iterations in ATE (m)

	Standing I	RoadCrossing III	RoadCrossing VII
1 iter	0.0649	0.1666	0.3157
2 iter	0.0315	0.0974	0.0658
3 iter	0.0327	0.0608	0.0660
Finetuned	0.0384	0.0631	0.0531

We then compare our method with others on the seven sequences from AirDOS-Shibuya in Tab. I and demonstrate that our method outperforms existing state-of-the-art VO algorithms. This benchmark covers much more challenging viewpoints and diverse motion patterns for articulated objects than our training data. The seven sequences are categorized into three levels of difficulty: most humans stand still in Standing Human with few of them moving around, Road

Crossing (Easy) contains multiple humans moving in and out of the camera’s view, and in Road Crossing (Hard) humans enter camera’s view abruptly. Besides VO methods, we also compare ours with SLAM methods that are able to handle dynamic scenes. DROID-SLAM [37] is a learning-based SLAM trained on TartanAir. AirDOS [38], VDO-SLAM [40] and DynaSLAM [41] are three feature-based SLAM methods targeting dynamic scenes. We provide the performance of AirDOS and ORB-SLAM [39] after masking the dynamic features during their ego-motion estimation. DeepVO [4], TartanVO and TrianFlow [42] are three learning-based VO methods not targeting dynamic scenes while CC [11] is an unsupervised VO resolving dynamic scenes through motion segmentation.

Our model achieves the best performance in all sequences among VO baselines and is competitive even among SLAM methods. DeepVO, TrianFlow and CC perform badly on AirDOS-Shibuya dataset because they are trained on KITTI only and not able to generalize. TartanVO performs better but it is still susceptible to the disturbance of dynamic objects. On RoadCrossing V as shown in Fig. 1, all VO baselines fail except ours. In hard sequences where there are more aggressive camera movements and abundant moving objects, ours outperforms dynamic SLAM methods such as AirDOS, VDO-SLAM and DynaSLAM by more than 80%. While DROID-SLAM remains competitive most time, it loses track of RoadCrossing V and VII as soon as a walking person occupies a large area in the image. Note that ours only takes 0.16 seconds per inference with 3 iterations but DROID-SLAM takes extra 4.8 seconds to optimize the trajectory. More qualitative results are in the supplemental material.

C. Performance on KITTI

We also evaluated our method against others on sequences from KITTI Odometry dataset [44] in Tab. II. Our method outperforms other VO baselines in 6 out of 8 dynamic sequences with an improvement of 27.7% on average against the second best method. DeepVO, TrianFlow and CC are trained on some of the sequences in KITTI while ours has not been finetuned on KITTI and is trained purely using synthetic

TABLE II: Results of ATE (m) on Dynamic Sequences from KITTI Odometry. Original sequences are trimmed into shorter ones that contain dynamic objects¹. DeepVO [4], TrianFlow [42] and CC [11] are trained on KITTI, while ours has not been finetuned on KITTI and is trained purely using synthetic data. Without backend optimization unlike SLAM, we achieve the best performance on 00, 02, 04, and competitive performance on the rest among all methods including SLAM.

		00	01	02	03	04	07	08	10
SLAM method	DROID-SLAM [37]	0.0148	49.193	0.1064	0.0119	0.0374	0.1939	0.9713	0.0368
	ORB-SLAM w/ mask [39]	0.0187	-	0.0796	0.1519	0.0198	0.2108	1.0479	0.0246
	DynaSLAM [41]	0.0138	-	0.1046	-	0.1450	0.3187	1.0559	0.0264
VO method	DeepVO [4]	(0.0206)	1.2896	(0.2975)	0.0783	0.0506	0.7262	(0.6547)	0.1042
	TrianFlow [42]	0.6966	(8.2127)	(1.8759)	1.6862	1.2950	1.5540	(3.8984)	0.2545
	CC [11]	0.0253	(0.3060)	(0.2559)	0.0505	<u>0.0337</u>	<u>0.6789</u>	(1.0411)	<u>(0.0346)</u>
	TartanVO [5]	0.0345	4.7080	<u>0.1049</u>	<u>0.2832</u>	0.0743	0.7108	<u>0.9776</u>	0.1024
	Ours	0.0126	<u>0.4081</u>	0.0594	0.0406	0.0180	0.6367	1.0344	0.0280

We use (·) to denote the sequence is in the training set of the corresponding method.

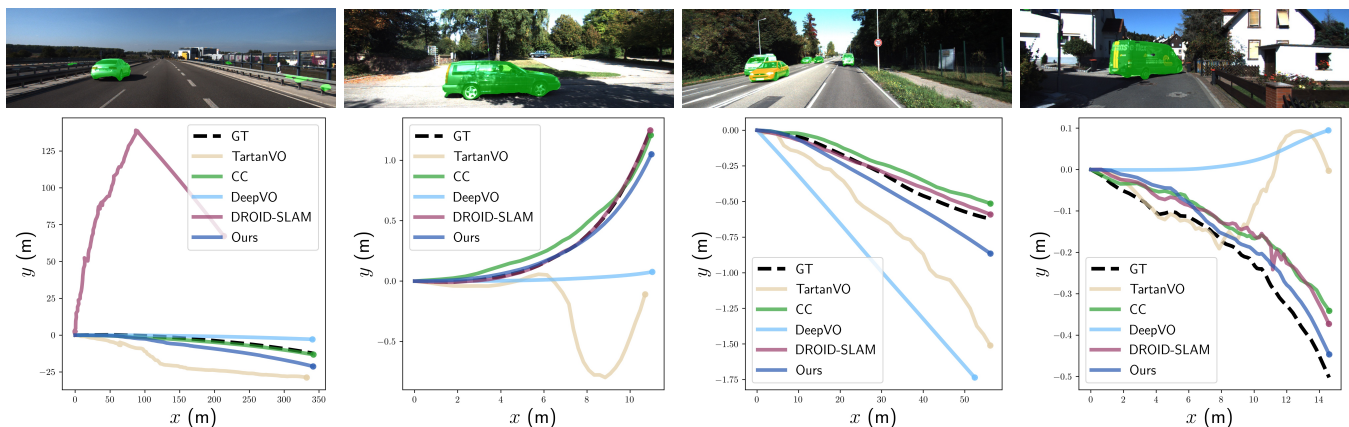


Fig. 4: Qualitative results on dynamic sequences in KITTI Odometry 01, 03, 04 and 10. The first row is our segmentation outputs of moving objects. The second row is the visualization after aligning the scales of trajectories with ground truth all at once. Ours produces precise odometry given large areas in the image being dynamic even among methods that are trained on KITTI. Note that the trajectories do not always reflect the ATE results due to alignment.

data. Moreover, we achieve the best ATE on 3 sequences among both VO and SLAM without any optimization. We provide qualitative results in Fig. 4 on four challenging sequences with fast-moving vehicles or dynamic objects occupying large areas in images. Note on sequence 01 which starts with a high-speed vehicle passing by, both ORB-SLAM and DynaSLAM fail to initialize, while DROID-SLAM loses track from the beginning. Even though CC uses 01 in its training set, ours gives only 0.1 higher ATE while 0.88 lower than the third best baseline. On sequence 10 when a huge van takes up significant areas in the center of the image, ours is the only VO that keeps track robustly.

D. Diagnostics

While we observe our method is robust to heavily dynamic scenes with as much as 70% dynamic objects in the image, it still fails when all foreground objects are moving, leaving textureless background only. This is most likely to happen when dynamic objects take up large areas in the image. For example, when testing on the test set of FlyingThings3D [30] where 80% of the image being dynamic, our method masks

¹Sequences listed are trimmed into lengths of 28, 133, 67, 31, 40, 136, 51 and 59 respectively which contain moving pedestrians, vehicles and cyclists.

almost the entire optical flow map as zeros, leading to the divergence of motion estimation and segmentation. Future work could hence consider incorporating dynamic object-awareness into the framework and utilizing dynamic cues instead of fully discarding them. Additionally, learning-based VO tends to overfit on simple translational movements such as in KITTI, which is resolved in our method by training on datasets with broad diversity, but our method gives worse performance when there is little or zero camera motion, caused by the bias in currently available datasets. One should consider training on zero-motion inputs in addition frame skipping.

V. CONCLUSION

In this paper, we propose a learning-based dynamic VO (DyTanVO) which can jointly refine the estimation of camera pose and segmentation of the dynamic objects. We demonstrate both ego-motion estimation and motion segmentation can converge quickly within time constraints for real-time applications. We evaluate our method on KITTI Odometry and AirDOS-Shibuya datasets, and demonstrate state-of-the-art performance in dynamic environments without finetuning nor optimization on the backend. Our work introduces new directions for dynamic visual SLAM algorithms.

REFERENCES

- [1] D. Scaramuzza and F. Fraundorfer, "Visual odometry [tutorial]," *IEEE robotics & automation magazine*, vol. 18, no. 4, pp. 80–92, 2011.
- [2] J. Engel, V. Koltun, and D. Cremers, "Direct sparse odometry," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 3, pp. 611–625, 2017.
- [3] C. Forster, M. Pizzoli, and D. Scaramuzza, "Svo: Fast semi-direct monocular visual odometry," in *2014 IEEE international conference on robotics and automation (ICRA)*, pp. 15–22, IEEE, 2014.
- [4] S. Wang, R. Clark, H. Wen, and N. Trigoni, "Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks," in *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 2043–2050, IEEE, 2017.
- [5] W. Wang, Y. Hu, and S. Scherer, "Tartanvo: A generalizable learning-based vo," *arXiv preprint arXiv:2011.00359*, 2020.
- [6] H. Zhou, B. Ummenhofer, and T. Brox, "Deeptam: Deep tracking and mapping," in *Proceedings of the European conference on computer vision (ECCV)*, pp. 822–838, 2018.
- [7] S. Li, X. Wang, Y. Cao, F. Xue, Z. Yan, and H. Zha, "Self-supervised deep visual odometry with online adaptation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6339–6348, 2020.
- [8] W. Xuan, R. Ren, S. Wu, and C. Chen, "Maskvo: Self-supervised visual odometry with a learnable dynamic mask," in *2022 IEEE/SICE International Symposium on System Integration (SII)*, pp. 225–231, IEEE, 2022.
- [9] U.-H. Kim, S.-H. Kim, and J.-H. Kim, "Simvodis++: Neural semantic visual odometry in dynamic environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4244–4251, 2022.
- [10] Z. Yin and J. Shi, "Geonet: Unsupervised learning of dense depth, optical flow and camera pose," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1983–1992, 2018.
- [11] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12240–12249, 2019.
- [12] H. Liu, G. Liu, G. Tian, S. Xin, and Z. Ji, "Visual slam based on dynamic object removal," in *2019 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pp. 596–601, IEEE, 2019.
- [13] B. Xu, W. Li, D. Tzoumanikas, M. Bloesch, A. Davison, and S. Leutenegger, "Mid-fusion: Octree-based object-level multi-instance dynamic slam," in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 5231–5237, IEEE, 2019.
- [14] S. Li and D. Lee, "RGB-D SLAM in dynamic environments using static point weighting," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2263–2270, 2017.
- [15] Y. Sun, M. Liu, and M. Q.-H. Meng, "Improving rgb-d slam in dynamic environments: A motion removal approach," *Robotics and Autonomous Systems*, vol. 89, pp. 110–122, 2017.
- [16] F. Fraundorfer and D. Scaramuzza, "Visual odometry: Part ii: Matching, robustness, optimization, and applications," *IEEE Robotics & Automation Magazine*, vol. 19, no. 2, pp. 78–90, 2012.
- [17] A. Valada, N. Radwan, and W. Burgard, "Deep auxiliary learning for visual localization and odometry," in *2018 IEEE international conference on robotics and automation (ICRA)*, pp. 6939–6946, IEEE, 2018.
- [18] G. Costante, M. Mancini, P. Valigi, and T. A. Ciarfuglia, "Exploring representation learning with cnns for frame-to-frame ego-motion estimation," *IEEE robotics and automation letters*, vol. 1, no. 1, pp. 18–25, 2015.
- [19] H. Zhan, C. S. Weerasekera, J.-W. Bian, and I. Reid, "Visual odometry revisited: What should be learnt?," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4203–4210, IEEE, 2020.
- [20] D.-H. Kim and J.-H. Kim, "Effective background model-based rgb-d dense visual odometry in a dynamic environment," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1565–1573, 2016.
- [21] S. Pillai and J. Leonard, "Monocular slam supported object recognition," *arXiv preprint arXiv:1506.01732*, 2015.
- [22] C. Yu, Z. Liu, X.-J. Liu, F. Xie, Y. Yang, Q. Wei, and Q. Fei, "Ds-slam: A semantic visual slam towards dynamic environments," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1168–1174, IEEE, 2018.
- [23] J. McCormac, A. Handa, A. Davison, and S. Leutenegger, "Semanticfusion: Dense 3d semantic mapping with convolutional neural networks," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4628–4635, IEEE, 2017.
- [24] L. An, X. Zhang, H. Gao, and Y. Liu, "Semantic segmentation-aided visual odometry for urban autonomous driving," *International Journal of Advanced Robotic Systems*, vol. 14, no. 5, p. 1729881417735667, 2017.
- [25] K.-N. Lianos, J. L. Schonberger, M. Pollefeys, and T. Sattler, "Vso: Visual semantic odometry," in *Proceedings of the European conference on computer vision (ECCV)*, pp. 234–250, 2018.
- [26] T. Zhang, H. Zhang, Y. Li, Y. Nakamura, and L. Zhang, "Flowfusion: Dynamic dense rgb-d slam based on optical flow," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7322–7328, IEEE, 2020.
- [27] W. Wang, D. Zhu, X. Wang, Y. Hu, Y. Qiu, C. Wang, Y. Hu, A. Kapoor, and S. Scherer, "Tartanair: A dataset to push the limits of visual slam," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4909–4916, IEEE, 2020.
- [28] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pp. 23–30, IEEE, 2017.
- [29] J. Tremblay, A. Prakash, D. Acuna, M. Brophy, V. Jampani, C. Anil, T. To, E. Cameracci, S. Bochoon, and S. Birchfield, "Training deep networks with synthetic data: Bridging the reality gap by domain randomization," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 969–977, 2018.
- [30] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox, "A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4040–4048, 2016.
- [31] D. Sun, X. Yang, M. Liu, and J. Kautz, "Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume," *CoRR*, vol. abs/1709.02371, 2017.
- [32] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241, Springer, 2015.
- [33] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015.
- [34] G. Yang and D. Ramanan, "Learning to segment rigid motions from two frames," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1266–1275, 2021.
- [35] G. Yang and D. Ramanan, "Upgrading optical flow to 3d scene flow through optical expansion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1334–1343, 2020.
- [36] G. French, A. Oliver, and T. Salimans, "Milking cowmask for semi-supervised image classification," *arXiv preprint arXiv:2003.12022*, 2020.
- [37] Z. Teed and J. Deng, "Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras," *Advances in Neural Information Processing Systems*, vol. 34, pp. 16558–16569, 2021.
- [38] Y. Qiu, C. Wang, W. Wang, M. Henein, and S. Scherer, "Airdos: Dynamic slam benefits from articulated objects," in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 8047–8053, IEEE, 2022.
- [39] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "Orb-slam: A versatile and accurate monocular slam system," *IEEE transactions on robotics*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [40] J. Zhang, M. Henein, R. Mahony, and V. Ila, "Vdo-slam: a visual dynamic object-aware slam system," *arXiv preprint arXiv:2005.11052*, 2020.
- [41] B. Bescos, J. M. Fácil, J. Civera, and J. Neira, "DynaSLAM: Tracking, mapping, and inpainting in dynamic scenes," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4076–4083, 2018.
- [42] W. Zhao, S. Liu, Y. Shu, and Y.-J. Liu, "Towards better generalization: Joint depth-pose learning without posenet," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9151–9161, 2020.
- [43] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., "Pytorch: An

imperative style, high-performance deep learning library,” *Advances in neural information processing systems*, vol. 32, 2019.

- [44] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “Vision meets robotics: The kitti dataset,” *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.