

Improved Real-Time Monocular SLAM Using Semantic Segmentation on Selective Frames

Jinkyu Lee, *Student Member, IEEE*, Muhyun Back, *Student Member, IEEE*,
Sung Soo Hwang*, *Member, IEEE*, and Il Yong Chun*, *Member, IEEE*

Abstract—Monocular simultaneous localization and mapping (SLAM) is emerging in advanced driver assistance systems and autonomous driving, because a single camera is cheap and easy to install. Conventional monocular SLAM has two major challenges leading inaccurate localization and mapping. First, it is challenging to estimate scales in localization and mapping. Second, conventional monocular SLAM uses inappropriate mapping factors such as dynamic objects and low-parallax areas in mapping. This paper proposes an improved real-time monocular SLAM that resolves the aforementioned challenges by efficiently using deep learning-based semantic segmentation. To achieve the real-time execution of the proposed method, we apply semantic segmentation only to downsampled keyframes in parallel with mapping processes. In addition, the proposed method corrects scales of camera poses and three-dimensional (3D) points, using estimated ground plane from road-labeled 3D points and the real camera height. The proposed method also removes inappropriate corner features labeled as moving objects and low parallax areas. Experiments with eight video sequences demonstrate that the proposed monocular SLAM system achieves significantly improved and comparable trajectory tracking accuracy, compared to existing state-of-the-art *monocular* and *stereo* SLAM systems, respectively. The proposed system can achieve real-time tracking on a standard CPU potentially with a standard GPU support, whereas existing segmentation-aided monocular SLAM does not.

Index Terms—Visual simultaneous localization and mapping (SLAM), Monocular SLAM, Keyframes, Deep semantic segmentation, Scale correction, Advanced driver assistance systems (ADAS), Autonomous driving.

I. INTRODUCTION

Simultaneous localization and mapping (SLAM) techniques have been evolving and widely applied to advanced driver assistance systems (ADAS) and autonomous driving systems.

S. S. Hwang is supported by Sabbatical Leave Grant from Handong Global University (HGU-2021). I. Y. Chun is supported in part by the National Research Foundation of Korea (NRF) grant NRF-2022R1F1A1074546 funded by the Ministry of Science of ICT (MSIT), the Institute of Information & communications Technology Planning & Evaluation (IITP) grant 2019-0-00421 funded by MSIT, the Institute for Basic Science grant (IBS) IBS-R015-D1, the Technology Innovation Program grant 20014967 funded by the Ministry of Trade, Industry & Energy (MOTIE), and the Korea Institute for Advancement of Technology grant P0022094 funded by the Ministry of Education-MOTIE.

*Corresponding authors

Jinkyu Lee and Muhyun Back are with the Department of Information and Communication Engineering, Handong Global University, Pohang 37554, South Korea (e-mails: 21931008@handong.edu; 21931005@handong.edu).

Sung Soo Hwang is with the School of Computer Science and Electrical Engineering, Handong Global University, Pohang 37554, South Korea (e-mail: sshwang@handong.edu).

I. Y. Chun is with the School of Electronic and Electrical Engineering and the Department of Artificial Intelligence, Sungkyunkwan University, and the Center for Neuroscience Imaging Research, IBS, Suwon, Gyeonggi-Do 16419, South Korea (e-mail: iyunchun@skku.edu).

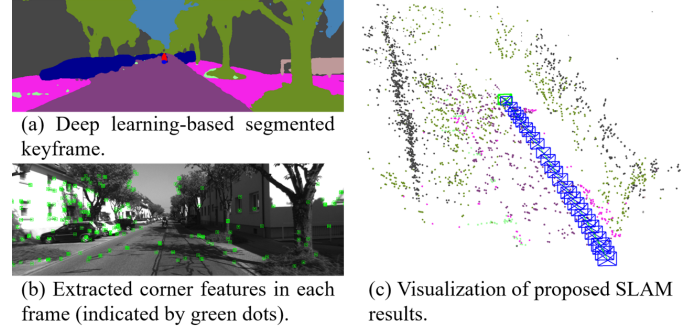


Fig. 1. Intermediate results of the proposed monocular SLAM framework. The proposed method uses segmented keyframes in (a), to choose appropriate corner features extracted in (b) and label reconstructed 3D points in (c). In (c), the colors of labeled 3D points correspond to those in the segmentation result (b). The proposed method uses road-labeled 3D points (colored with purple in (c)) for scale estimation.

While SLAM approaches using light detection and ranging (LiDAR) sensors are accurate, the cost of LiDAR sensors is high and they have not been widely used in commercial products. Visual SLAM systems that use camera(s) are a popular alternative to LiDAR-based SLAM. Monocular SLAM systems that use a single camera are attractive as they are cheap and easy to install. Monocular SLAM was initially suggested with filter-based approaches [1]–[4]. The filter-based methods are computationally inefficient, since both localization and mapping run on every frame [5]. To resolve the issue of filter-based methods, keyframe-based approaches [6]–[8] (see other references in [5]) run the mapping process only on selective frames, called keyframes, while the localization process estimates a camera pose in every frame. The keyframe-based SLAM improved the localization accuracy and computational efficiency of filter-based methods [9], and became the *de facto* standard in monocular SLAM [5].

Conventional monocular SLAM has major challenges in scale estimation and mapping. There are two scale estimation challenges in monocular SLAM, namely, scale ambiguity and scale drift [10]. First, one cannot recover camera poses and three-dimensional (3D) points with absolute scale. Second, errors are accumulated by following the SLAM procedures, i.e., feature extraction, feature matching, camera pose estimation, and 3D structure reconstruction. In monocular SLAM, the scale ambiguity is the unique problem and scale drifts become more severe due to the scale ambiguity. If the scale issues exist, monocular SLAM can mis-localize a current camera (or vehicle) position and generate 3D points in “off-the-grid” locations. The other major challenge is inaccurate mapping

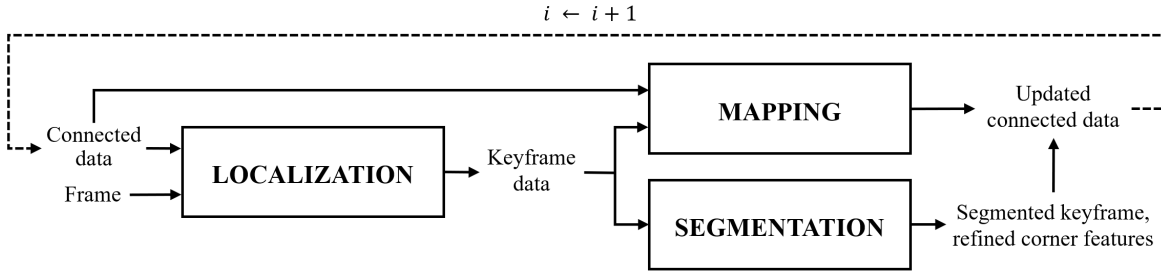


Fig. 2. Overview of the proposed monocular SLAM system between previous and current keyframe time points, indexed by $i - 1$ and i , respectively. The keyframe data consists of current keyframe and its corner features and camera pose, and previous keyframe and its corrected camera pose. These are denoted by $\mathbf{k}^{(i)}$, $\mathbf{f}^{(i)}$, $\mathbf{P}^{(i)}$, $\mathbf{k}^{(i-1)}$, $\hat{\mathbf{P}}^{(i-1)}$, respectively; see notations in Section II-A. The connected data consists of refined corner features, corrected camera pose, segmentation result and corrected 3D points of connected keyframes. These are denoted by $\hat{\mathbf{f}}^{(i')}$, $\hat{\mathbf{P}}^{(i')}$, $\hat{\mathbf{k}}^{(i')}$, $\hat{\mathbf{x}}^{(i')}$ for $i' \in \mathcal{C}^{(i)}$, respectively; see notations in Section II-A. The symbol colors correspond to those in Fig. 3.

from two-dimensional (2D) points to 3D points. The mapping process of monocular SLAM consists of the triangulation of 2D point pairs obtained from two adjacent frames collected from different time points [11], where we refer to a point in time as a time point. For accurate mapping, static scenes and sufficient parallax are required. In practice, however, some objects such as cars, pedestrians, etc. moves in a scene, or some objects are distant from a camera with low-parallax.

A. Related methods for overcoming limitations of monocular SLAM

This section reviews related works that attempted to resolve limitations of monocular SLAM.

To resolve the scale ambiguity issue, many existing methods use prior knowledge, e.g., camera height [10], [12]–[17], object size [18]–[20], and pedestrian height [19], and estimate the absolute scale in monocular SLAM. To reduce scale drift, [21] proposed a loop detection method that recognizes revisiting prior locations. However, many driving trajectories do not complete a loop. [10] uses a camera height to estimate the absolute scale and correct scale drifts. The method estimates the ground plane from a fixed region of interest (ROI) that corresponds to road area(s). The estimation performance is limited if an ROI is blocked by obstacles such as preceding vehicles. To resolve this limitation, this work uses road areas labeled by semantic segmentation.

To improve mapping accuracy, researchers combined monocular SLAM with deep learning-based semantic segmentation. [22], [23] use semantic segmentation to probabilistically classify static and dynamic areas. The Mask-SLAM method excludes dynamic objects and sky area, by using a mask produced by semantic segmentation [24]. [25] combines semantic segmentation with a moving consistency-check method to deal with dynamic objects. These approaches showed excellent performance even for scenes with dynamic objects. These methods, however, use a few segmented labels including dynamic objects and sky areas. One may use more labels, such as road areas, to estimate ground planes. More importantly, deep learning-based semantic segmentation has high computational costs, and applying it to each frame could terminate monocular SLAM. This limits the practical use of this powerful combination in real-time applications such as autonomous driving. Different from existing methods [22]–[25],

this work applies semantic segmentation only to keyframes to achieve real-time monocular SLAM.

B. Contributions and organization of the paper

This paper proposes an improved real-time monocular SLAM framework that solves the aforementioned limitations of monocular SLAM, by using deep learning-based semantic segmentation in sophisticated ways. The main contributions of this paper are as follows:

- To resolve the time complexity problem of segmentation-based monocular SLAM, we apply deep learning-based semantic segmentation *only* to keyframes. To further improve the efficiency, segmentation process is performed in a different thread in parallel with mapping processes.
- To overcome the scale estimation challenges, we estimate the absolute scale using segmented road labels and the actual camera height, and apply it to correct scales of camera poses and 3D points.
- To reduce inaccurate mappings, we filter out, i.e., refine, corner features of both moving objects, such as cars, pedestrians, etc., and low-parallax areas, that may generate inappropriate 3D points. To discriminate low-parallax areas, we use the Kanade–Lucas–Tomasi (KLT) tracker [26] that is widely used in the optical flow, and select corner features with small displacement in the sky, terrain, and building areas.
- The proposed monocular SLAM framework using semantic segmentation simultaneously resolves the aforementioned scale estimation and mapping challenges.

Fig. 1 shows some intermediate results of the proposed monocular SLAM framework.

The remainder of this paper is organized as follows. Section II describes the proposed efficient integration approach of deep learning-based semantic segmentation with monocular SLAM, and the proposed scale correction and inaccurate mapping reduction methods by labeling 3D points and 2D corner features labeled based on segmented keyframes, respectively. Section III reports experiments where the proposed monocular SLAM system achieves significantly improved and comparable trajectory tracking accuracy, compared to existing state-of-the-art *monocular* SLAMs and *stereo* SLAM, respectively. Furthermore, Section III shows that the proposed system

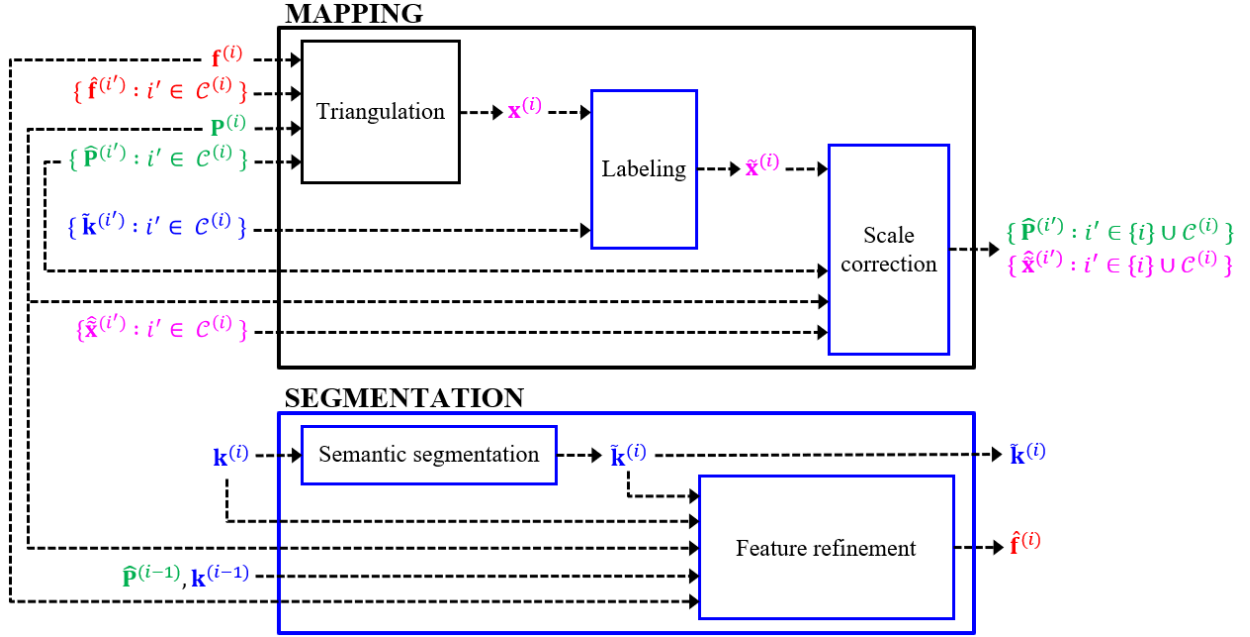


Fig. 3. Proposed combination of keyframe-based monocular SLAM and semantic segmentation. Blue boxes indicate modified or added (sub-)modules. Red symbols indicate corner features, green symbols indicate camera pose matrix, blue symbols indicate keyframes, and magenta symbols indicate 3D point map. The input and output data sets are listed on left and right, respectively.

significantly reduces tracking time per frame compared to existing integration of segmentation into monocular SLAM [24], achieving real-time processing on a standard CPU potentially with a standard GPU support. Section IV concludes the study and includes some future works.

II. PROPOSED METHOD

The proposed monocular SLAM system consists of three modules, localization, mapping, and segmentation, where each module runs in different threads. We modify the backbone of the conventional keyframe-based monocular SLAM method [8], efficiently integrating segmentation. We summarize processes of each module between previous and current keyframe time point. The localization module selects a keyframe when both mapping and segmentation modules complete processing on a previous keyframe; similar to [8], it provides corner features and camera pose of a current keyframe (extracting corner features of each frame and estimating its camera pose with 3D points of connected keyframes). The mapping module generates new 3D points by using triangulation between current corner features and connected corner features [8], and assigns segmented labels from connected segmented keyframes to new 3D points. It then estimates a ground plane using only road-labeled 3D points and correct a scale of camera poses and 3D points using estimated ground plane and real camera height. While the mapping module provides scale-corrected camera poses and 3D points, the segmentation module performs deep learning-based segmentation to a downsampled keyframe and refines keyframe's corner features by removing moving objects and low-parallax areas using segmentation result. Fig. 2 overviews the proposed monocular SLAM system.

Section II-A describes notations for data in the proposed method. Section II-B describes how we integrate deep

learning-based segmentation with keyframe-based monocular SLAM. Using segmented keyframe results, Sections II-C and II-D describe details of scale correction in 3D mapping and removal of factors in a current keyframe that potentially cause inappropriate mapping, respectively.

A. Preliminaries

This section summarizes notations for data in the proposed method:

- The superscript $(\cdot)^{(i)}$ denotes the i th keyframes index. (i) indicates the current keyframe time point.
- $(\hat{\cdot})$ denotes segmented image or points with semantic labels. $(\tilde{\cdot})$ denotes that input is refined or scale-corrected.
- $\mathbf{k}^{(i)}$ denotes the i th keyframe. $\tilde{\mathbf{k}}$ denotes segmented keyframe. Fig. 1(a) shows a segmented keyframe.
- $\mathbf{f}^{(i)}$ denotes a 2D image with ORB corner features [27] extracted from $\mathbf{k}^{(i)}$ via the localization module. Fig. 1(b) shows extracted corner features. $\hat{\mathbf{f}}$ denotes refined corner features by removing factors that potentially cause inappropriate mapping, via the segmentation module.
- $\mathbf{P}^{(i)}$ denotes 3×4 normalized camera pose matrix [11] of $\mathbf{k}^{(i)}$. Normalized camera pose matrix is defined by $\mathbf{P} \triangleq [\mathbf{R}, \mathbf{t}]$, where \mathbf{R} and \mathbf{t} denote a 3D rotation matrix and translation vector, respectively. These are calculated by the localization module. $\hat{\mathbf{P}}$ denotes scale-corrected normalized camera pose matrix, where its translation vector is calculated by scale-corrected coordinate of the camera center. A camera center coordinate is calculated by $\mathbf{c} = -\mathbf{R}^{-1}\mathbf{t}$ [11], where $(\cdot)^{-1}$ indicates the inverse of a matrix. $\hat{\mathbf{c}}$ denotes scale-corrected camera center coordinate calculated from $\hat{\mathbf{P}}$.

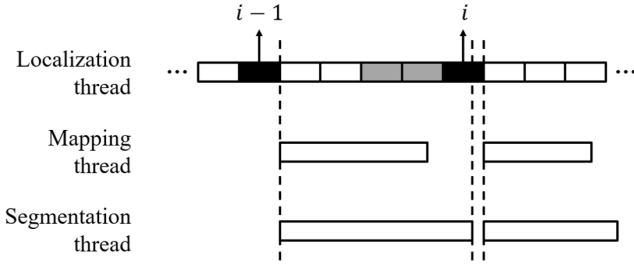


Fig. 4. Proposed keyframe selection strategy and its timing. In localization thread, the length of each box indicates the processing time for each frame. In mapping and segmentation thread, the length of each box indicates the processing time for each keyframe. Gray and black box denotes keyframe candidate and selected keyframe, respectively. We choose a keyframe only after both mapping and segmentation processes are completed, although localization thread flags frame(s) as keyframe candidate(s). Each candidate is determined at the end of each frame localization.

- $\mathbf{x}^{(i)}$ denotes a 3D map with 3D points generated from $\mathbf{f}^{(i)}$ via the mapping module. $\hat{\mathbf{x}}$ denotes a scale-corrected 3D map, where a scale of 3D point coordinate are corrected. $\hat{\tilde{\mathbf{x}}}$ denotes a scale-corrected 3D map with semantic labels. Fig. 1(c) illustrates 3D points with semantic labels.
- $\mathcal{C}^{(i)}$ denotes a set of indices for connected keyframes with the i th keyframe, in previous keyframe time points. If two keyframes generate 3D points at equal to or greater than 15 same coordinates, they are called *connected* [8]. Specifically, $\mathcal{C}^{(i)} := \{i-1, \dots, i-C^{(i)}\}$, where $C^{(i)}$ is the number of connected keyframes with $\mathbf{k}^{(i)}$.

B. Proposed efficient integration of semantic segmentation with monocular SLAM

To efficiently integrate semantic segmentation with monocular SLAM, we apply semantic segmentation only to keyframes. Applying semantic segmentation to monocular SLAM will improve the monocular SLAM accuracy by preventing generating inappropriate 3D points [22]–[25] and assigning semantic labels to 3D points for better ground-plane estimation. However, applying semantic segmentation to every frame is computational demanding and unsuitable for real-time monocular SLAM. 3D points are generated from the mapping process that uses only keyframes, so segmenting only keyframes is sufficient for processing 3D points in monocular SLAM.

In addition, we perform the mapping and segmentation process in two different threads. In monocular SLAM, 3D points are generated by triangulating matched features from two keyframes [11]. Two keyframes have matched features, so segmenting only one of two keyframes may be sufficient for subsequent 3D point processing. The proposed integration approach performs the mapping process with connected segmented keyframes and related data in one thread, in parallel with performing the segmentation process with a current keyframe and related data in the other thread. Fig. 3 illustrates the parallel processing. We describe details of each thread processing below.

In the mapping thread, we first generate a 3D point map $\mathbf{x}^{(i)}$ by triangulation between current corner features $\mathbf{f}^{(i)}$ and each of connected (refined) corner features $\{\hat{\mathbf{f}}^{(i')} : i' \in \mathcal{C}^{(i)}\}$,

using a current camera matrix $\mathbf{P}^{(i)}$ and each of connected (scale-corrected) camera matrices $\{\hat{\mathbf{P}}^{(i')} : i' \in \mathcal{C}^{(i)}\}$. We then assign semantic labels in $\{\tilde{\mathbf{k}}^{(i')} : i' \in \mathcal{C}^{(i)}\}$ to generated $\mathbf{x}^{(i)}$ and obtain $\tilde{\mathbf{x}}^{(i)}$. Finally, we correct a scale of keyframe’s camera center coordinates computed by each of $\{\mathbf{P}^{(i)}, \hat{\mathbf{P}}^{(i')} : i' \in \mathcal{C}^{(i)}\}$ and that of 3D point coordinates in each of $\{\tilde{\mathbf{x}}^{(i)}, \hat{\tilde{\mathbf{x}}}^{(i')} : i' \in \mathcal{C}^{(i)}\}$. While the mapping thread performs aforementioned processes, the segmentation thread first segments a current keyframe $\mathbf{k}^{(i)}$, and obtains $\tilde{\mathbf{k}}^{(i)}$. We then refine a current corner feature $\mathbf{f}^{(i)}$, using current data $\{\mathbf{P}^{(i)}, \mathbf{k}^{(i)}, \tilde{\mathbf{k}}^{(i)}\}$ and previous data $\{\hat{\mathbf{P}}^{(i-1)}, \mathbf{k}^{(i-1)}\}$, and obtain $\hat{\mathbf{f}}^{(i)}$. Fig. 3 summarizes the above processes. We will describe details of the scale correction and corner feature refinement process later in Section II-C and II-D, respectively.

Note that we choose a keyframe only after both the mapping and segmentation processes are completed. If the localization thread generates keyframe candidates in the middle of mapping or segmentation process, we disregard them and choose the one generated right after the mapping and segmentation processes are completed as a keyframe. (It is likely that subsequent frames after a keyframe candidate are also keyframe candidates.) Fig. 4 summarizes this keyframe selection strategy. Specifically, we select $\{\mathbf{f}^{(i)}, \mathbf{P}^{(i)}, \mathbf{k}^{(i)}\}$ from the localization module and begin the mapping and segmentation processes, only after $\{\hat{\mathbf{P}}^{(i-1)}, \hat{\tilde{\mathbf{x}}}^{(i-1)}\}$ and $\{\tilde{\mathbf{k}}^{(i-1)}, \hat{\mathbf{f}}^{(i-1)}\}$ are obtained by the mapping and segmentation thread in the previous keyframe time point, respectively.

Lastly, we downsample each 2D keyframe to reduce computational complexity of semantic segmentation (and upsample segmented keyframes back to the original resolution). Considering the above keyframe selection strategy and observations that semantic segmentation is generally more computationally expensive than the mapping module, the segmentation computation time critically affects the decision timing of a new keyframe. The slower semantic segmentation, the longer time gap between keyframes. This leads to fewer matched corner features between two keyframe time points and fewer generated 3D points via triangulation. The localization module projects 3D points onto each frame and estimates camera pose from projected 3D points and extracted corner features [8]. If the number of 3D points is small, then pose estimation accuracy reduces and eventually monocular SLAM performance significantly degrades. Reducing the keyframe image size/resolution can decrease the computational costs of semantic segmentation methods using convolutional neural network (CNN), at a cost of segmentation accuracy. Our experiments disregarding segmentation CNN computation costs demonstrate that although CNN segmentation accuracy decreases by reduced image size, the proposed monocular SLAM gives comparable performance with that using the original image size. See details later in Section III-B.

C. Scale correction

At each current keyframe time point, we correct the scale of keyframe’s camera center coordinate and each 3D point coordinate. We first estimate ground plane parameters $\{\mathbf{n}^{(i)}, d^{(i)}\}$ using only *road* labeled 3D points in $\{\tilde{\mathbf{x}}^{(i)}, \hat{\tilde{\mathbf{x}}}^{(i')} : i' \in \mathcal{C}^{(i)}\}$.

A ground plane is parameterized with a normal vector $\mathbf{n}^{(i)}$ ($\|\mathbf{n}\|_2 = 1$) and the perpendicular distance between a ground plane and the origin, $d^{(i)}$, using the plane equation in 3D $\langle \mathbf{n}, (x, y, z) \rangle = 0$ [28], where $\langle \cdot, \cdot \rangle$ denotes the inner product between two vectors and (x, y, z) is a coordinate of any 3D point on a plane. For accurate ground plane estimation, we use a robust parameter estimation method, random sample consensus (RANSAC) [29] that uses only inliers of road-labeled 3D points (we normalized \mathbf{n} at each RANSAC iteration). Different from existing methods that use unlabeled 3D points in ROI for RANSAC, e.g., [10], we use 3D points with road label.

We then compute a scaling factor [12] using *virtual* camera height and a *real* camera height h_{real} that is measured in advance. We estimate a *virtual* camera height $h^{(i)}$ with the estimated ground plane parameters $\{\mathbf{n}^{(i)}, d^{(i)}\}$ above. A *virtual* camera height $h^{(i)}$ is the perpendicular distance between estimated ground plane and a current keyframe's camera center coordinate $\mathbf{c}^{(i)}$, and calculated by $h^{(i)} = |\langle \mathbf{n}^{(i)}, \mathbf{c}^{(i)} \rangle + d^{(i)}|$. We compute a scaling factor $s^{(i)}$ using h_{real} and $h^{(i)}$:

$$s^{(i)} = \frac{h_{\text{real}}}{h^{(i)}}. \quad (1)$$

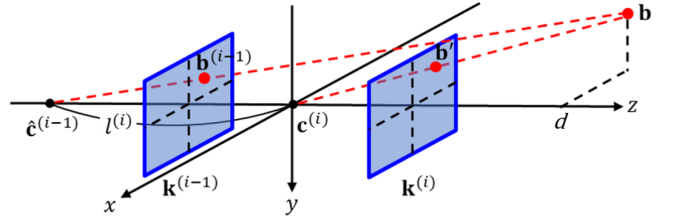
Intuitively speaking, $s^{(i)}$ above quantifies the scale drift at the current keyframe time point.

Finally, we adjust coordinates of keyframe's camera center computed by each of $\{\mathbf{P}^{(i)}, \hat{\mathbf{P}}^{(i')} : i' \in \mathcal{C}^{(i)}\}$ and that of each 3D point in $\{\tilde{\mathbf{x}}^{(i)}, \hat{\tilde{\mathbf{x}}}^{(i')} : i' \in \mathcal{C}^{(i)}\}$, using the scaling factor $s^{(i)}$ in (1). Applying $s^{(i)}$ follows the scale correction processes in [10]. We perform scale correction when the scaling factor satisfies $0.001 < |s^{(i)} - 1| < 0.2$, to avoid frequent or intensive correction. We adjust camera center coordinates of keyframes, by scaling the distance between $\mathbf{c}^{(i)}$ and a camera center coordinate calculated from each of $\{\hat{\mathbf{P}}^{(i')} : i' \in \mathcal{C}^{(i)}\}$ with $s^{(i)}$. Likewise, we adjust 3D point coordinates by scaling the distance between $\mathbf{c}^{(i)}$ and each 3D point coordinate in $\{\hat{\tilde{\mathbf{x}}}^{(i')} : i' \in \mathcal{C}^{(i)}\}$ with $s^{(i)}$.

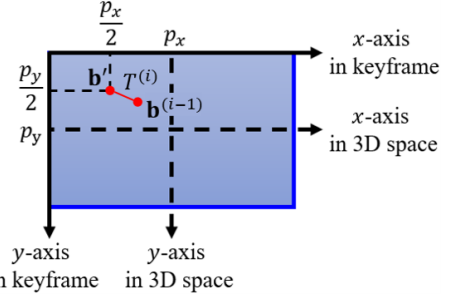
In the first several keyframes, we do not perform the scale correction since road-labeled 3D points are insufficient. When the number of road-labeled 3D points become large or equal than 50, we estimate $h^{(i)}$ by averaging the absolute differences between $\mathbf{c}^{(i)}$'s y -coordinate and y -coordinates of road-labeled 3D points in $\{\tilde{\mathbf{x}}^{(i')} : i' \in \{i\} \cup \mathcal{C}^{(i)}\}$, and correct scales of $\{\tilde{\mathbf{x}}^{(i')}, \mathbf{P}^{(i')} : i' \in \{i\} \cup \mathcal{C}^{(i)}\}$ with $s^{(i)}$ in (1). (RANSAC uses a user-specific threshold in determining inlier samples and this threshold depends on scales of samples [12], but such 3D points are not yet scale-corrected in this keyframe time point.) Subsequent keyframes will use RANSAC-based ground plane estimation for the scale correction process, as described earlier.

D. Corner feature refinement by removing inappropriate mapping factors

At each current keyframe time point, we refine current corner features $\mathbf{f}^{(i)}$ by removing corner features of moving objects and low-parallax areas, using current data $\{\mathbf{P}^{(i)}, \mathbf{k}^{(i)}, \tilde{\mathbf{k}}^{(i)}\}$ and previous data $\{\hat{\mathbf{P}}^{(i-1)}, \mathbf{k}^{(i-1)}\}$. Using a current segmented keyframe $\tilde{\mathbf{k}}^{(i)}$, we first remove corner features in areas labeled as movable objects, such as car, person, and bicycle, etc., similar to the existing methods [22]–[24]. We next remove



(a) 3D pin-hole camera projection on current and previous keyframes.



(b) $T^{(i)}$ is the distance between \mathbf{b}' and $\mathbf{b}^{(i-1)}$ on 2D keyframe plane. The principal point coordinate (p_x, p_y) is defined on 2D keyframe plane.

Fig. 5. Geometrical illustrations of the proposed adaptive $T^{(i)}$ selection using 3D pin-hole projection model on current and previous keyframes.

corner features labeled as sky, terrain, and building – referred to as backgrounds – in low-parallax areas. For each corner feature in backgrounds, we compute the displacement between a current keyframe $\mathbf{k}^{(i)}$ and a previous keyframe $\mathbf{k}^{(i-1)}$, using KLT tracker [26]. Corner features with the displacement less than some threshold are considered as low-parallax areas.

When a 3D point is projected on a current keyframe and a previous keyframe, the displacement between projected 2D points changes depending on the distance between two keyframes. We adaptively select a threshold $T^{(i)}$ at each current keyframe time point, using the 3D projection model on two keyframes (see below) and estimated distance between current and previous keyframes.

This paragraph describes the 3D pin-hole camera projection model on current and previous keyframes, and its notations. A 3D point with coordinate \mathbf{b} projects onto $\mathbf{k}^{(i)}$ - and $\mathbf{k}^{(i-1)}$ -keyframe 2D planes (towards their camera centers with coordinates $\mathbf{c}^{(i)}$ and $\hat{\mathbf{c}}^{(i-1)}$, respectively). The mapped points on current and previous keyframes are denoted by \mathbf{b}' and $\mathbf{b}^{(i-1)}$, respectively. The distance between $\mathbf{c}^{(i)}$ and $\hat{\mathbf{c}}^{(i-1)}$ along the z -direction is denoted by $l^{(i)}$. Fig. 5(a) illustrates this. In setting or estimating \mathbf{b} , \mathbf{b}' , and $\mathbf{b}^{(i-1)}$, we use camera intrinsic parameters, focal length f and principal point coordinate (p_x, p_y) (obtained by camera calibration in advance).

The proposed adaptive $T^{(i)}$ selection consists of four steps. First, we set \mathbf{b}' using principal point coordinate:

$$\mathbf{b}' = \begin{bmatrix} \frac{p_x}{2} \\ \frac{p_y}{2} \end{bmatrix}, \quad (2)$$

where we observed that in general, low-parallax areas are located near this coordinate; see its location on a 2D keyframe plane in Fig. 5(b). Second, we set \mathbf{b} as follows: we set the z -coordinate of \mathbf{b} as d by assuming that inappropriate mapping

occurs in objects at a distance of d m or more and calculate x - and y -coordinates of \mathbf{b} by

$$\mathbf{b} = \begin{bmatrix} -\frac{dp_x}{2f} \\ -\frac{dp_y}{2f} \\ d \end{bmatrix}. \quad (3)$$

From this perspective, we refer to d as the low-parallax feature determination parameter. We obtained the result in (3) by back-projecting \mathbf{b}' in (2) to \mathbf{b} using the perspective projection equation [11]:

$$\lambda \begin{bmatrix} \mathbf{b}' \\ 1 \end{bmatrix} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \mathbf{b},$$

where we set $\lambda = d$ to express \mathbf{b}' in the form of homogeneous coordinate. Here, the 3×3 matrix is the camera intrinsic matrix. Third, we estimate $\mathbf{b}^{(i-1)}$ using $l^{(i)}$:

$$\mathbf{b}^{(i-1)} = \begin{bmatrix} \frac{p_x(l^{(i)}+d/2)}{l^{(i)}+d} \\ \frac{p_y(l^{(i)}+d/2)}{l^{(i)}+d} \end{bmatrix}. \quad (4)$$

We obtained the result in (4) by mapping \mathbf{b} in (3) with the z -coordinate increased by $l^{(i)}$ onto $\mathbf{k}^{(i-1)}$ -keyframe 2D plane (using the perspective projection):

$$\lambda \begin{bmatrix} \mathbf{b}^{(i-1)} \\ 1 \end{bmatrix} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} -\frac{dp_x}{2f} \\ -\frac{dp_y}{2f} \\ l^{(i)} + d \end{bmatrix},$$

where we set $\lambda = l^{(i)} + d$ to express $\mathbf{b}^{(i-1)}$ in the form of homogeneous coordinate. Finally, we select a threshold $T^{(i)}$ by calculating the distance between \mathbf{b}' in (2) and $\mathbf{b}^{(i-1)}$ in (4):

$$T^{(i)} = \|\mathbf{b}' - \mathbf{b}^{(i-1)}\|_2 = \frac{l^{(i)}}{2(l^{(i)} + d)} \sqrt{p_x^2 + p_y^2}. \quad (5)$$

This is illustrated in Fig. 5(b). Objects more than d m away from the $\mathbf{c}^{(i)}$ -point may have the displacement less than $T^{(i)}$. We remove the corresponding corner features with displacement less than $T^{(i)}$. We determine an appropriate low-parallax feature determination parameter d by comparing performances of proposed monocular SLAM with different d values (see details later in III-C1).

III. RESULTS AND DISCUSSION

We compared the proposed monocular SLAM with four variations of ORB-SLAM2 [30]: 1) monocular setup without loop closing (LC), 2) monocular setup with LC, 3) stereo setup without LC, and 4) stereo setup with LC. We additionally compared the proposed method with Mask-SLAM [24], a representative monocular SLAM method that uses semantic segmentation. The ORB-SLAM2 stereo setup uses two cameras and resolves the scale ambiguity & drift [30] and inaccurate mapping issues [31]. LC can further improve the estimation accuracy of camera poses and 3D point coordinates in ORB-SLAM2, if it detects a loop. It can be applied to both monocular and stereo setups. Mask-SLAM applies semantic segmentation to all frames to exclude dynamic objects. We

mainly focus on comparing the trajectory tracking accuracy between different visual SLAMs that do *not* experience tracking failure, where tracking failure may additionally need relocalization [32]–[34].

In measuring the trajectory tracking accuracy, we calculated absolute trajectory error (ATE) [35], root-mean-squared-error between predicted camera center coordinates of keyframes and their ground-truth in meters (m) (if ground-truth exists). Throughout experiments, we did not use the manual re-alignment process [8, Sec. VIII-E] that realigns an estimated trajectory to the ground-truth, since this is impractical.

A. Experimental setup

We first determined an appropriate downsampling factor for keyframe segmentation (see the last paragraph in Section II-B), and compared the proposed method using the determined downsampling factor with the five state-of-the-art visual SLAM methods described above. We used either a GPU or CPU for semantic segmentation of keyframe images, and ran remaining modules on a CPU. We used 3.1 GHz Intel Core i9 with 128 GB RAM and 1607 MHz NVIDIA GeForce GTX 1080 with 8GB RAM, for CPU processing and GPU processing, respectively.

In the localization module, we extracted 3,000 ORB corner features [27] for each input frame.

1) *Comparisons of segmentation and tracking accuracies and segmentation running time, with different keyframe (spatial) resolutions:* We compared the accuracy of semantic segmentation and the trajectory tracking and running time of semantic segmentation of the proposed monocular SLAM system with different keyframe resolutions. To generate keyframes with different resolutions, we downsampled the width and height of original keyframes by factors of 4/3, 2, or 4 using bicubic downsampling. We upsampled segmented keyframes with nearest neighbor interpolation (using an upsampling factor identical to corresponding downsampling factor). In the proposed method, long segmentation running time(s) can jeopardize camera pose estimation in localization (due to limited 3D points), so we disregarded segmentation computation costs in measuring trajectory tracking accuracy.

For keyframe segmentation with a GPU support, we used the advanced NVIDIA semantic segmentation network [36].¹ This network was pre-trained with the KITTI semantic segmentation dataset [39] that is different from the KITTI odometry dataset [40]. The proposed monocular SLAM method uses four classes, road, movable objects, background, and others, where we categorized predictions of the segmentation network as follows: road (“road”), movable objects (“person”, “rider”, “car”, “truck”, “bus”, “train”, “motorcycle”, and “bicycle”), background (“building”, “terrain”, and “sky”), and others.

For CPU-based keyframe segmentation, we used ERFNet [41] that can achieve high computational efficiency by using residual connections and factorized convolutions. Because ERFNet is designed to use input with the specific aspect ratio

¹This method uses video prediction based data synthesis [37] to increase the training data size, and it took the first place in the KITTI semantic segmentation benchmark [38].

(width : height = 2 : 1), we adjusted the size of keyframe images to fit the ratio. (We only tested the factor 1 and 4/3 for ERFNet.) This network was also trained with the KITTI semantic segmentation dataset [38], with hyperparameters given in [41] except for 200 epochs and a batch size of 6. We categorized predictions of ERFNet into the aforementioned four classes used in the NVIDIA network.

In measuring the trajectory tracking accuracy and segmentation running time, we used the sequence 00 of the KITTI odometry dataset [40]. In measuring the segmentation accuracy, we used the Cityscapes dataset [42] that consists of 3,475 images driving in urban environments, since the KITTI-00 dataset does not have ground-truth segmentation annotations. We evaluated the semantic segmentation accuracy by mean intersection over union (IoU) of the road, movable objects, and background labels (others were not used in the proposed method).

We determined an appropriate spatial downsampling factor for keyframe segmentation considering trade-off between trajectory tracking accuracy and segmentation running time.

2) *Comparisons of proposed monocular SLAM with different sub-module combinations:* We ran the proposed monocular SLAM system with different low-parallax feature determination parameters $d = 200, 250, 300$ and chose an appropriate one by considering both the trajectory tracking accuracy and tracking success rate. We used the chosen parameter throughout the experiments.

In addition, we ran an ablation study by removing the scale correction and feature refining sub-modules, respectively. We measured the trajectory tracking accuracy from each setup.

Throughout the experiments in this section, we used the GPU-based NVIDIA network with a determined keyframe resolution in Section III-A1, and the sequence 00 of the KITTI odometry dataset [40].

3) *Comparisons between different visual SLAM methods:* We ran the proposed monocular SLAM considering segmentation computation times with a determined keyframe spatial resolution via the trade-off investigation experiment in Section III-B. We tested both segmentation processing methods, the NVIDIA network on a GPU and ERFNet on a CPU. Segmentation running time includes computational costs of downsampling, semantic segmentation, upsampling, and data transfer times between CPU and GPU (if GPU is used). We segmented all downsampled frames in advance; when a keyframe is selected, we added the corresponding segmentation running time to the segmentation thread.

For comparing the trajectory tracking accuracy with different visual SLAM methods, we used the sequences 00, 02, 03, 04, 05, 07, and 08 of the KITTI odometry dataset [40]. The KITTI odometry dataset consists of sequence sets, where each sequence set is obtained by driving in a different residential area. The original image size is 1241×376 and the camera frame-rate is 10 Hz. The sequences 00, 02, 05, and 07 have a loop or loops in their trajectories. The sequences 03, 04, and 08 do not have loop(s), so we only ran ORB-SLAM2 without LC (for both monocular and stereo setups). To consider dynamic environments with moving objects, we used the KITTI 02 and 04 sequences.

In addition, we used a highway sequence provided by StradVision Inc. [43] to consider environments with both moving objects and distant objects. The StradVision sequence includes monocular video of 900 frames where many cars drive on the highway (i.e., there exist many moving objects) surrounded by mountains (i.e., low-parallax area). The StradVision dataset was obtained with imaging resolution 1280×720 and camera frame-rate 30 Hz. In experiments with the StradVision dataset, we compared the proposed method using the NVIDIA network, with monocular ORB-SLAM2. The sequence does not have any loops, so we only ran monocular ORB-SLAM2 without LC.

In addition to the quantitative analyses with ATEs, we qualitatively compared trajectory tracking accuracy in the KITTI odometry dataset by comparing estimated trajectories with their ground-truths. Since the StradVision sequence does not have ground-truth trajectory, we aligned trajectories with the map provided by Google Maps for qualitative comparison. (In monocular ORB-SLAM2 without LC, we also aligned a scale of the estimated trajectory.)

To study the feasibility of running semantic segmentation with monocular SLAM in real-time, we calculated the tracking time per frame in milliseconds (ms). We compared the proposed efficient combination framework of monocular SLAM and semantic segmentation, with the existing integration, Mask-SLAM. In measuring the tracking time per frame, we used all the aforementioned KITTI sequences and calculated its average and standard deviation values.

B. Comparisons of segmentation and tracking accuracies and segmentation running time, with different keyframe (spatial) resolutions

1) *Segmentation accuracy comparisons with different keyframe resolutions:* For GPU-based semantic segmentation with the NVIDIA network, Table I shows that when keyframes are downsampled by a factor of 4/3 or 2, segmentation accuracy (slightly) decreases, compared to the case with the original keyframe resolution. When keyframes are downsampled by a factor of 4, the segmentation accuracy significantly decreases, compared to the case with the original keyframe resolution. Specifically, the segmentation accuracy (in IoU) decreases by 3.85%, 9.60%, and 7.04% in road, movable objects, and background labels, respectively. Comparing the results of the NVIDIA network and ERFNet without downsampling in Table I shows that ERFNet running on CPU slightly degrades the segmentation accuracy of the NVIDIA network on GPU. The trend applies to comparisons with NVIDIA segmentation network (on GPU) using a downsampling factor up to 2; see the NVIDIA network results in Table. Comparing the downsampling factor 4/3 results between the NVIDIA network and ERFNet in Table I shows that ERFNet on CPU significantly degrades the segmentation accuracy, compared to the NVIDIA network on GPU.

2) *Tracking accuracy comparisons with different keyframe resolutions:* For GPU-based semantic segmentation with NVIDIA network, Table I shows that although keyframes are downsampled by a factor of 4/3 or 2, the trajectory tracking

TABLE I
SEGMENTATION ACCURACY, RUNNING TIME AND TRAJECTORY TRACKING ACCURACY WITH DIFFERENT PROCESSING METHODS AND DIFFERENT KEYFRAME DOWNSAMPLING FACTORS IN PROPOSED MONOCULAR SLAM (KITTI-00 DATASET)

Segmentation methods	Downsampling factor (spatial resolution)	Segmentation accuracy (IoU)			Segmentation running time (ms)		ATE (m)
		Road	Movable obj.	Background	Mean	Std.	
NVIDIA network on GPU	1 (1241×376)	0.934	0.865	0.867	479.05	6.72	6.721
	4/3 (930×282)	0.926	0.863	0.863	288.44	5.90	7.755
	2 (620×188)	0.918	0.859	0.855	129.44	4.13	6.784
	4 (310×94)	0.898	0.782	0.806	61.66	4.94	8.681
ERFNet on CPU	≈1 (960×480)	0.912	0.627	0.825	238.90	13.66	8.031
	4/3 (720×360)	0.877	0.551	0.782	131.99	22.20	11.932

accuracy is similar to the case with the original keyframe resolution. When keyframes are downsampled by a factor of 4, trajectory tracking accuracy significantly decreases, compared to the case without downsampling. Comparing the results of the NVIDIA network and ERFNet without downsampling in Table I shows that using ERFNet running on CPU decreases the tracking accuracy of the proposed method using the NVIDIA network on GPU. The trend applies to comparisons with NVIDIA network segmentation (on GPU) using a downsampling factor up to 2; see the NVIDIA network results in Table I. Comparing the downsampling factor 4/3 results between the NVIDIA network and ERFNet in Table I shows that ERFNet on CPU significantly degrades the trajectory tracking accuracy.

3) *Segmentation running time comparisons with different keyframe resolutions:* Table I shows that segmentation running time approximately linearly decreases, as the keyframe downsampling factor increases, regardless of network and processing unit choices. This is natural, as the computational complexity of segmentation network scales with the size of input images. Note that CPU-based ERFNet without downsampling runs segmentation faster than the GPU-based NVIDIA network using the downsampling factor 4/3. Compare the corresponding results in Table I.

4) *Real-time processing feasibility of the proposed system with different keyframe resolutions in segmentation:* When segmentation computation times are considered (see Section III-A3), we observed with the KITTI-00 sequence that the proposed monocular SLAM system runs continuously with the GPU-based NVIDIA networks using the downsampling factors 2 and 4, and CPU-based ERFNets using no downsampling and the downsampling factor 4/3; with the downsampling factors 1 and 4/3 in GPU-based NVIDIA net, the proposed system frequently stops running. The reason is that long segmentation running time can significantly delay keyframe selection and limit 3D point generation, and eventually, cause the entire system failure; see further details in Section II-B. We observed that if the NVIDIA network runs on a standard CPU, the proposed system eventually stops running even with the keyframe downsampling factor 4 for segmentation.

To achieve real-time execution of the proposed monocular visual SLAM system on the test datasets, we decide to spatially downsample keyframes by the factor 2 in the GPU-

based NVIDIA network and to use no downsampling in CPU-based ERFNet for semantic segmentation. For the proposed system with NVIDIA segmentation network running on a standard GPU, the downsampling factor 2 for segmentation shows a reasonable trade-off between tracking accuracy and segmentation running time. This is supported by the empirical observations in Section III-B. The keyframe downsampling factor 2 for segmentation 1) slightly degrades the trajectory tracking accuracy (only by 0.94%), but 2) significantly reduces the segmentation running time (by a factor of ≈ 4), so proposed monocular SLAM runs continuously. For the proposed system with ERFNet running a standard CPU, we decide not to use downsampling for keyframes in segmentation, because even with the downsampling factor 4/3, the trajectory tracking accuracy significantly degrades.

One may have a different favorable keyframe downsampling factor for other datasets. We suggest to choose a keyframe downsampling factor (for segmentation) that achieves segmentation accuracy higher than 0.9 IoU for road labels.

C. Comparisons of proposed monocular SLAM with different module combinations

1) *Comparisons of proposed monocular SLAM with different low-parallax feature determination parameters:* Table II shows that as the low-parallax feature determination parameter d increases, the number of filtered corner features and the trajectory tracking accuracy decrease. When d is set to 200, the trajectory tracking accuracy is slightly better compared to the $d = 250$ case. However, the proposed system frequently stops running, as shown by the low tracking success rate in Table II, where the tracking success rate denotes the ratio of number of frames in which localization is successfully executed, to total number of frames. This is because too many corner features are removed, when $d = 200$. Therefore we set the low-parallax feature determination parameter to 250 for the proposed corner feature refinement.

2) *Ablation study of proposed monocular SLAM:* Comparing Table III(b) and III(d) shows that proposed monocular SLAM with the scale correction sub-module dramatically improves the trajectory tracking accuracy compared to monocular ORB-SLAM2 without LC. Comparing Table III(c) and III(d) shows that proposed monocular SLAM with the feature refinement sub-module significantly improves the trajectory

TABLE II
AVERAGED NUMBER OF FILTERED CORNER FEATURES, TRAJECTORY TRACKING ACCURACY, AND TRACKING SUCCESS RATE WITH DIFFERENT LOW-PARALLAX DETERMINATION FEATURE DETERMINATION PARAMETERS IN PROPOSED MONOCULAR SLAM (KITTI-00 DATASET)

Low-parallax feature determination parameter	Averaged number of filtered corner features	ATE (m)	Tracking success rate (%)
200	291.24	6.409	42.06
250	96.75	6.797	100
300	36.26	8.459	100

TABLE III
ABLATION STUDY OF SCALE CORRECTION AND FEATURE REFINEMENT SUB-MODULES IN PROPOSED MONOCULAR SLAM (KITTI-00 DATASET)

	(a) Proposed mono. SLAM w/ both sub-modules	(b) Proposed mono. SLAM w/ only scale correction	(c) Proposed mono. SLAM w/ only feature refinement	(d) Mono. ORB-SLAM2 w/o LC
ATE (m)	6.797	9.253	212.616	231.701

tracking accuracy compared to monocular ORB-SLAM2 without LC. In particular, the scale correction sub-module plays an important role in improving proposed monocular SLAM, by resolving the scale ambiguity issue that is a major challenge in monocular SLAM. The proposed monocular SLAM system can maximize the trajectory tracking accuracy when it uses both the scale correction and the feature refinement sub-modules that require image segmentation process.

D. Comparisons between different visual SLAM methods

1) *Trajectory tracking accuracy with different visual SLAM methods:* In all the sequences of the KITTI odometry dataset, regardless of the loop(s) existence, Tables IV(a)–IV(d) & IV(g)–IV(h) and comparing their trajectories in Figs. 6–7 show that the proposed monocular SLAM with either NVIDIA network on GPU or ERFNet on CPU significantly improves trajectory tracking accuracy compared to monocular ORB-SLAM2 regardless of using LC, and Mask-SLAM. ERFNet slightly degrades the tracking accuracy of the proposed method, compared to the NVIDIA network. This is due to the degradation of segmentation accuracy discussed in Section III-B.

The proposed monocular visual SLAM system resolves the scale ambiguity by correcting scales using segmented road labels, and achieves significantly better trajectory tracking accuracy. On the other hand, the monocular ORB-SLAM2 methods suffer from scale ambiguity regardless of LC, their estimated trajectories are off the scale from ground-truth trajectories, and ATE values significantly increase. Similar to monocular ORB-SLAM2, Mask-SLAM suffers from the scale ambiguity and drift issues, giving poor results in Tables II(g)–II(h) and Fig. 8.

Regardless of the loop(s) existence in the KITTI odometry sequences, Tables IV(a)–IV(b), and IV(e) and comparing their trajectories in Figs. 6–7 demonstrate that the trajectory tracking accuracy of the proposed method with either GPU-based NVIDIA network or CPU-based ERFNet is comparable to that

of stereo ORB-SLAM2 without LC. The results imply that using only a single camera, the proposed method can resolve scale ambiguity & drift and inaccurate mapping issues with comparable performance to stereo ORB-SLAM2 without LC.

In the sequences with loop(s), stereo ORB-SLAM2 with LC gave slightly better trajectory tracking accuracy compared to the proposed method. See Tables IV(a)–IV(b) and IV(f) and compare their results in Fig. 6. In both monocular and stereo ORB-SLAM2s, the LC scheme improves trajectory tracking accuracy by correcting accumulated localization errors, when a vehicle revisits a previously mapped area (i.e, a loop is created). Compare results between Tables IV(c) and IV(d), and those between and Tables IV(e) and IV(f). We conjecture by considering the results in Tables IV(a)–IV(b) and Table IV(e) that if LC is additionally used, the proposed monocular SLAM system achieves comparable or better trajectory tracking accuracy, compared to stereo ORB-SLAM2 with LC.

In the StradVision sequence (that has no loops), Fig. 9 shows that the proposed method significantly improves the trajectory tracking accuracy compared to monocular ORB-SLAM2 without LC. Whereas the proposed method consistently followed a road route, monocular ORB-SLAM2 without LC gave deviated vehicle localization at the 535th frame, and was eventually terminated at the 749th frame. The monocular ORB-SLAM2 system using no LC does not remove moving objects and low-parallax areas in generating 3D points at its mapping module. Such inappropriate 3D points could significantly degrade camera pose estimation performance, and eventually may terminate the entire system. The proposed SLAM system overcomes that challenge by corner feature refinements in the segmentation thread; see Section II-D.

2) *Comparisons of real-time execution feasibility of different monocular SLAMs using semantic segmentation:* Table V compares tracking time per frame between the proposed monocular SLAM system and Mask-SLAM, given either NVIDIA network on a standard GPU or ERFNet on a standard CPU (see CPU and GPU specs in Section III-A). Regardless of segmentation methods, the averaged tracking time per frame of the proposed system is below the inverse of the camera frame-rate, 100 ms, implying that the system can run in real-time [30] on a standard CPU and/or GPU. The proposed efficient integration of semantic segmentation with monocular SLAM achieves real-time execution by 1) applying semantic segmentation to keyframes and 2) sophisticatedly considering localization and mapping timings; see details in Section II-B. Since the averaged tracking time per frame of the proposed method is below 33 ms given imaging resolution 1241×376 , we conjecture that the proposed method runs in real-time with some 30 Hz cameras with an imaging resolution lower than or equal to 1241×376 . Different from the proposed method, Table V implies that Mask-SLAM is infeasible for real-time running on a standard CPU or GPU. This aligns with the claim in [24, Sec. 5].

The proposed monocular SLAM – with either GPU-based segmentation using the NVIDIA network and the keyframe downsampling factor 2 or CPU-based segmentation using ERFNet – did not experience system termination/failure in all the eight sequences. One might operate the proposed monoc-

TABLE IV
ATE VALUES (M) OF DIFFERENT VISUAL SLAM METHODS IN DIFFERENT KITTI ODOMETRY SEQUENCES

Sequence		(a) Proposed mono. SLAM + NVIDIA [†]	(b) Proposed mono. SLAM + ERFNet [‡]	(c) Monocular ORB-SLAM2 without LC	(d) Monocular ORB-SLAM2 with LC	(e) Stereo ORB-SLAM2 without LC	(f) Stereo ORB-SLAM2 with LC	(g) Mask-SLAM + NVIDIA [†]	(h) Mask-SLAM + ERFNet [‡]
With loop(s)	00	6.797	8.451	231.701	199.658	8.515	6.869	229.908	231.473
	02	13.059	14.874	550.845	421.711	14.330	13.437	523.694	516.256
	05	2.888	3.763	170.815	143.368	2.949	1.567	153.925	154.101
	07	2.354	3.464	80.490	67.158	2.856	1.057	71.778	72.629
No loops	03	1.306	2.087	140.736	-	1.387	-	191.354	193.940
	04	1.271	1.730	200.135	-	1.008	-	191.348	191.578
	08	12.973	13.428	196.803	-	12.765	-	195.471	203.622

[†] The NVIDIA semantic-segmentation network [36] ran on a standard GPU with keyframes downsampled by factor 2.

[‡] ERFNet segmentation [41] ran on a standard CPU with non-downsampled keyframes.

TABLE V
TRACKING TIME PER FRAME WITH DIFFERENT MONOCULAR SLAM METHODS USING SEGMENTATION (KITTI SEQUENCES IN TABLE IV)

Monocular SLAM methods using segmentation	Tracking time per frame (ms)	
	Mean	Std.
Proposed mono. SLAM + NVIDIA [†]	28.52	7.77
Proposed mono. SLAM + ERFNet [‡]	29.74	9.95
Mask-SLAM + NVIDIA [†]	145.51	9.96
Mask-SLAM + ERFNet [‡]	270.63	9.30

[†] The NVIDIA semantic-segmentation network [36] ran on a standard GPU with keyframes downsampled by factor 2.

[‡] ERFNet segmentation [41] ran on a standard CPU with non-downsampled keyframes.

TABLE VI
PROCESSING TIME FOR EACH (SUB-)MODULES (KITTI-00 DATASET)

Modules	Sub-modules	Processing time (ms)	
		Mean	Std.
Localization (per frame)	Total	27.87	7.93
Mapping (per keyframe)	Triangulation	130.10	43.01
	Scale correction	14.99	32.45
	Total	146.55	73.64
Segmentation (per keyframe)	Semantic segmentation [†]	129.44	4.13
	Feature refinement	70.31	30.60
	Total	201.74	31.27

[†] The NVIDIA semantic-segmentation network [36] ran on a standard GPU with keyframes downsampled by factor 2.

ular SLAM system with the NVIDIA segmentation network for high-quality tracking in real-time, using some GPU-based vehicle board, e.g., NVIDIA Jetson [44]. If a standard GPU is unavailable, the proposed system with ERFNet may run in real-time on a standard CPU, but at little expense of tracking accuracy.

We note that the processing time of the localization module determines the real-time execution feasibility of the proposed method, because all modules run in different threads in parallel and the localization module is the only module that runs for every frames. In Table VI, the averaged processing times of the mapping and the segmentation modules are greater than 100 ms, but they run in parallel only for every *keyframe*.

IV. CONCLUSION

Monocular SLAM methods use only a single camera and thus, have major limitations in scale estimation and mapping accuracy. The proposed monocular SLAM system resolves the limitations by efficiently using deep learning-based semantic segmentation: it applies a state-of-the-art semantic segmentation only to (downsampled) keyframes in parallel with mapping processes. The proposed method performs scale correction using road-labeled 3D points and removes features labeled by dynamic objects and low parallax areas in 3D point generations. Regardless of the existence of loops in video sequences, proposed monocular SLAM achieves significantly better trajectory tracking accuracy compared to state-of-the-art monocular SLAM, monocular ORB-SLAM2 [30]. In video sequences without loops, proposed monocular SLAM achieves comparable trajectory tracking accuracy compared to state-of-the-art stereo SLAM, stereo ORB-SLAM2 [30]. Compared with an existing segmentation-aided monocular SLAM method, Mask-SLAM [24], the proposed system achieves significantly improved trajectory tracking accuracy and reduced tracking time per frame; it can run in real time on a standard CPU potentially with a standard GPU. We believe that the proposed monocular SLAM system would be particularly useful in ADAS and autonomous driving systems, by using only a single camera but having comparable performances to stereo visual SLAM systems. The proposed efficient integration of semantic segmentation with monocular SLAM could be applicable to a visual SLAM integration with other deep-learning methods needing high computation costs, such as depth estimation.

There are a number of avenues for future work. First, we will additionally run LC in an additional thread, to further improve trajectory tracking accuracy. Second, we additionally incorporate a relocalization technique (e.g., [32]–[34]) into the proposed method to improve the robustness against tracking failure. Third, we plan to use camera poses, 3D point coordinates, etc. to refine segmentation results, motivated by [45]. On the image sensor side, the future work is modifying the proposed framework for fish-eye or omnidirectional camera.

REFERENCES

- [1] A. J. Davison, I. D. Reid, N. D. Molton, and O. Stasse, "MonoSLAM: Real-time single camera SLAM," *IEEE Transactions on Pattern Analysis*

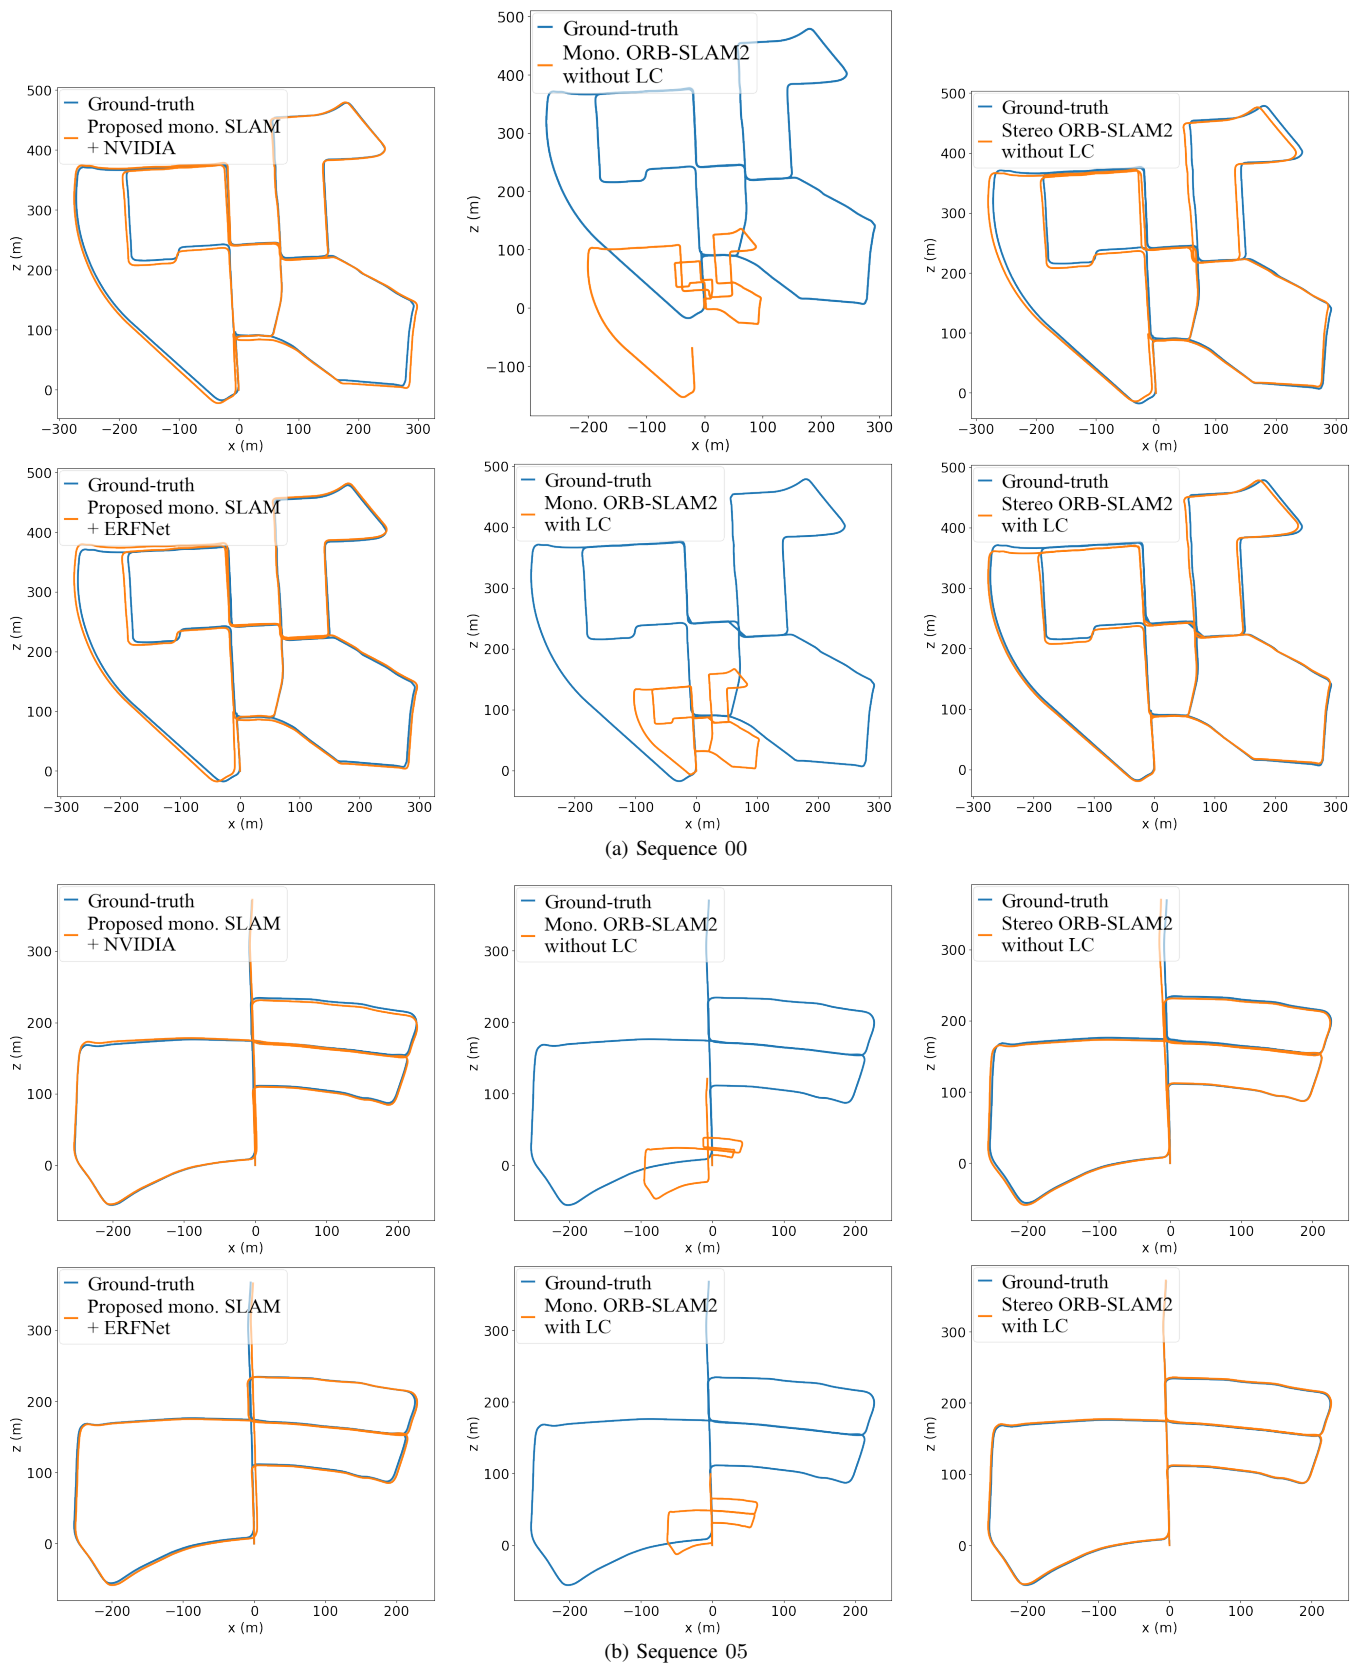


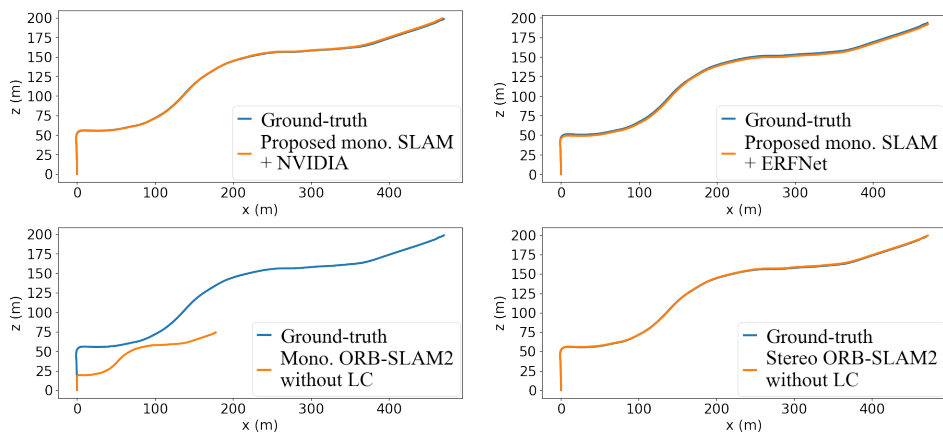
Fig. 6. Qualitative tracking comparisons in the KITTI odometry dataset. Sequence 00 and 05 have a loop or loops in trajectories.

and Machine Intelligence, vol. 29, no. 6, pp. 1052–1067, June 2007.

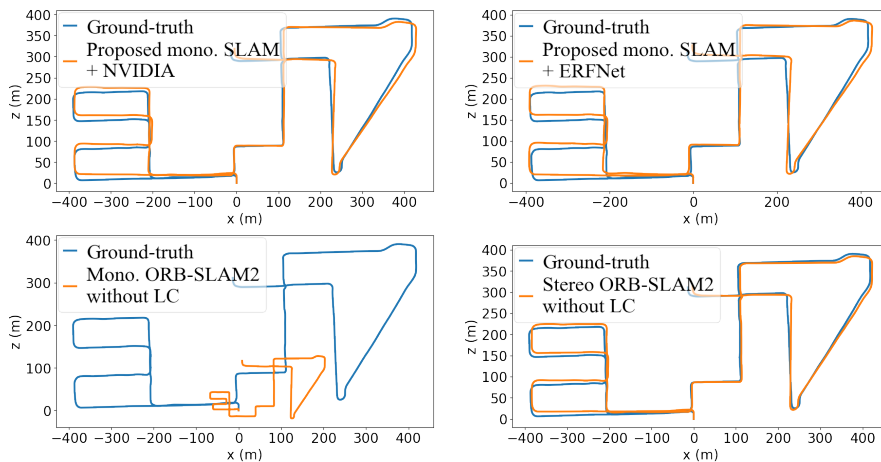
[2] J. Civera, A. J. Davison, and J. M. Montiel, “Inverse depth parametrization for monocular SLAM,” *IEEE Transactions on Robotics*, vol. 24,

no. 5, pp. 932–945, October 2008.

[3] A. Chiuso, P. Favaro, H. Jin, and S. Soatto, “Structure from motion causally integrated over time,” *IEEE Transactions on Pattern Analysis*

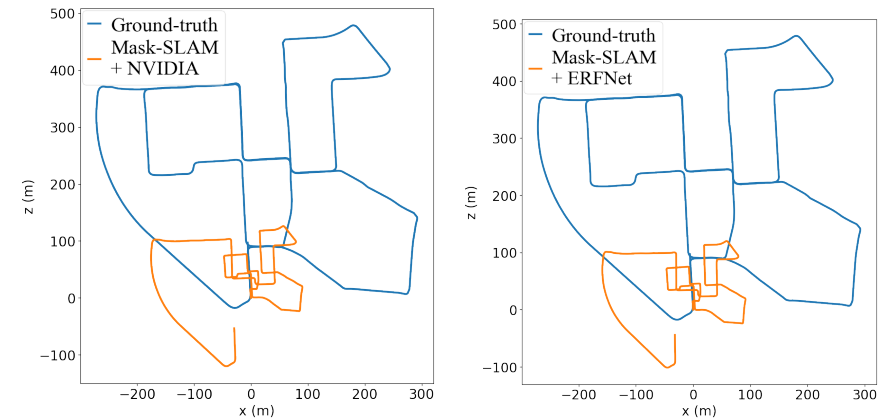


(a) Sequence 03

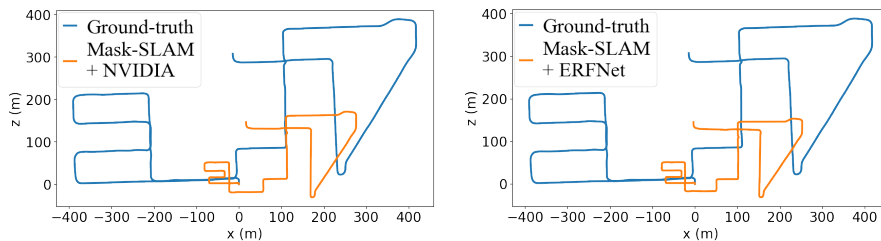


(b) Sequence 08

Fig. 7. Qualitative tracking comparisons in the KITTI odometry dataset. Sequence 03 and 08 have no loop(s) in trajectories.



(a) Sequence 00



(b) Sequence 08

Fig. 8. Qualitative tracking results of Mask-SLAM with different segmentation networks in the KITTI 00 and 08 sequences.

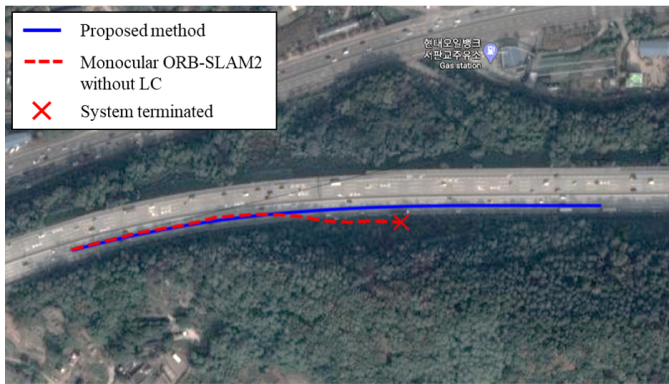


Fig. 9. Comparison of the proposed method (with the NVIDIA network on a standard GPU), and monocular ORB-SLAM2 without LC in the StradVision sequence. The blue solid and red dashed lines denote estimated trajectory by the proposed method and monocular ORB-SLAM2 without LC, respectively. The estimated trajectories are aligned with the map in Google Maps. We observed a very similar tracking performance between the proposed systems using the GPU-based NVIDIA net and and CPU-based ERFNet.

- and Machine Intelligence*, vol. 24, no. 4, pp. 523–535, April 2002.
- [4] E. Eade and T. Drummond, “Scalable monocular SLAM,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, New York, NY, June 2006, pp. 469–476.
- [5] G. Younes, D. Asmar, E. Shammas, and J. Zelek, “Keyframe-based monocular SLAM: design, survey, and future directions,” *Robotics and Autonomous Systems*, vol. 98, pp. 67–88, December 2017.
- [6] E. Mouragnon, M. Lhuillier, M. Dhome, F. Dekeyser, and P. Sayd, “Real time localization and 3d reconstruction,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, vol. 1, New York, NY, June 2006, pp. 363–370.
- [7] G. Klein and D. Murray, “Parallel tracking and mapping for small ar workspaces,” in *Proc. IEEE and ACM International Symposium on Mixed and Augmented Reality*, Nara, Japan, November 2007, pp. 225–234.
- [8] 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, August 2015.
- [9] H. Strasdat, J. M. Montiel, and A. J. Davison, “Visual SLAM: why filter?” *Image and Vision Computing*, vol. 30, no. 2, pp. 65–77, February 2012.
- [10] D. Zhou, Y. Dai, and H. Li, “Ground plane based absolute scale estimation for monocular visual odometry,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 2, pp. 791–802, May 2019.
- [11] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, March 2004.
- [12] J. Gräter, T. Schwarze, and M. Lauer, “Robust scale estimation for monocular visual odometry using structure from motion and vanishing points,” in *Proc. IEEE Intelligent Vehicles Symposium*, Seoul, Korea, June 2015, pp. 475–480.
- [13] S. Lovegrove, A. J. Davison, and J. Ibanez-Guzmán, “Accurate visual odometry from a rear parking camera,” in *Proc. IEEE Intelligent Vehicles Symposium*, Baden-Baden, Germany, June 2011, pp. 788–793.
- [14] M. H. Mirabdollah and B. Mertsching, “Fast techniques for monocular visual odometry,” in *Proc. German Conference on Pattern Recognition*, Aachen, Germany, October 2015, pp. 297–307.
- [15] F. I. Pereira, G. Ilha, J. Luft, M. Negreiros, and A. Susin, “Monocular visual odometry with cyclic estimation,” in *Proc. SIBGRAPI Conference on Graphics, Patterns and Images*, Niteroi, Brazil, October 2017, pp. 1–6.
- [16] D. Scaramuzza, F. Fraundorfer, M. Pollefeys, and R. Siegwart, “Absolute scale in structure from motion from a single vehicle mounted camera by exploiting nonholonomic constraints,” in *Proc. IEEE International Conference on Computer Vision*, Kyoto, Japan, September 2009, pp. 1413–1419.
- [17] S. Yang and S. Scherer, “CubeSLAM: Monocular 3-D object SLAM,” *IEEE Transactions on Robotics*, vol. 35, no. 4, pp. 925–938, August 2019.
- [18] T. Botterill, S. Mills, and R. Green, “Bag-of-words-driven, single-camera simultaneous localization and mapping,” *Journal of Field Robotics*, vol. 28, no. 2, pp. 204–226, March 2011.
- [19] —, “Correcting scale drift by object recognition in single-camera SLAM,” *IEEE Transactions on Cybernetics*, vol. 43, no. 6, pp. 1767–1780, December 2012.
- [20] S. Song, M. Chandraker, and C. C. Guest, “High accuracy monocular SFM and scale correction for autonomous driving,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 4, pp. 730–743, April 2016.
- [21] H. Strasdat, J. Montiel, and A. J. Davison, “Scale drift-aware large scale monocular SLAM,” *Robotics: Science and Systems*, vol. 2, no. 3, p. 7, June 2010.
- [22] 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, pp. 1–11, October 2017.
- [23] N. Brasch, A. Bozic, J. Lallemand, and F. Tombari, “Semantic monocular SLAM for highly dynamic environments,” in *Proc. IEEE/RISJ International Conference on Intelligent Robots and Systems*, Madrid, Spain, October 2018, pp. 393–400.
- [24] M. Kaneko, K. Iwami, T. Ogawa, T. Yamasaki, and K. Aizawa, “Mask-SLAM: Robust feature-based monocular SLAM by masking using semantic segmentation,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops*, Salt Lake City, UT, June 2018, pp. 258–266.
- [25] 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 *Proc. IEEE/RISJ International Conference on Intelligent Robots and Systems*, Madrid, Spain, October 2018, pp. 1168–1174.
- [26] S. Baker and I. Matthews, “Lucas-Kanade 20 years on: A unifying framework,” *International journal of computer vision*, vol. 56, no. 3, pp. 221–255, February 2004.
- [27] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, “ORB: An efficient alternative to SIFT or SURF,” in *Proc. IEEE International Conference on Computer Vision*, Barcelona, Spain, November 2011, pp. 2564–2571.
- [28] H. Anton and C. Rorres, *Elementary Linear Algebra*. John Wiley & Sons, September 2013.
- [29] M. A. Fischler and R. C. Bolles, “Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography,” *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, June 1981.
- [30] 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*, vol. 33, no. 5, pp. 1255–1262, June 2017.
- [31] D. Scaramuzza and F. Fraundorfer, “Visual odometry: Part 1: The first 30 years and fundamentals,” *IEEE robotics & automation magazine*, vol. 18, no. 4, pp. 80–92, December 2011.
- [32] E. Eade and T. Drummond, “Unified loop closing and recovery for real time monocular SLAM,” in *Proc. British Machine Vision Conference*, vol. 13, Leeds, UK, June 2008, p. 136.
- [33] B. Williams, G. Klein, and I. Reid, “Automatic relocalization and loop closing for real-time monocular SLAM,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 9, pp. 1699–1712, September 2011.
- [34] R. Mur-Artal and J. D. Tardós, “Fast relocalisation and loop closing in keyframe-based SLAM,” in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, Hong Kong, China, May 2014, pp. 846–853.
- [35] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, “A benchmark for the evaluation of RGB-D SLAM systems,” in *Proc. IEEE/RISJ International Conference on Intelligent Robots and Systems*, Vilamoura-Algarve, Portugal, October 2012, pp. 573–580.
- [36] Y. Zhu, K. Sapra, F. A. Reda, K. J. Shih, S. Newsam, A. Tao, and B. Catanzaro, “Improving semantic segmentation via video propagation and label relaxation,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, Long Beach, CA, June 2019, pp. 8856–8865.
- [37] F. A. Reda, G. Liu, K. J. Shih, R. Kirby, J. Barker, D. Tarjan, A. Tao, and B. Catanzaro, “SDC-Net: Video prediction using spatially-displaced convolution,” in *Proc. IEEE European Conference on Computer Vision*, Munich, Germany, September 2018, pp. 718–733.
- [38] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “The KITTI semantic segmentation benchmark,” Available from http://www.cvlibs.net/datasets/kitti/eval_semseg.php?benchmark=semantics2015, 2018.
- [39] H. Alhajjja, S. Mustikovela, L. Mescheder, A. Geiger, and C. Rother, “Augmented reality meets computer vision: Efficient data generation for urban driving scenes,” *International Journal of Computer Vision*, March 2018.

- [40] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, Providence, RI, June 2012, pp. 3354–3361.
- [41] E. Romera, J. M. Alvarez, L. M. Bergasa, and R. Arroyo, "ERFNet: Efficient residual factorized convnet for real-time semantic segmentation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 1, pp. 263–272, Jan 2018.
- [42] 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 *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, June 2016, pp. 3213–3223.
- [43] StradVision, "<https://stradvision.com/>," 2021.
- [44] NVIDIA Jetson Modules, "<https://developer.nvidia.com/embedded/jetson-modules>," 2021.
- [45] K. Wang, Y. Lin, L. Wang, L. Han, M. Hua, X. Wang, S. Lian, and B. Huang, "A unified framework for mutual improvement of SLAM and semantic segmentation," in *Proc. IEEE International Conference on Robotics and Automation*, Montreal, QC, Canada, May 2019, pp. 5224–5230.



Sung Soo Hwang was born in Busan, South Korea in 1983. He received B.S. degree in Electrical Engineering and Computer Science from Handong Global University, Pohang, South Korea in 2008, and M.S and Ph. D degrees in Korea Advanced Institute of Science and Technology, Daejeon, South Korea in 2010 and 2015, respectively.

He is an Associate Professor with School of Computer Science and Electrical Engineering, Handong Global University, Pohang, South Korea. His research interests include simultaneous localization and mapping system, autonomous driving, and image-based 3D modeling.



Il Yong Chun (S'14–M'16) is a tenure-track Assistant Professor of EEE & AI at Sungkyunkwan University (SKKU), Suwon, South Korea. He received B.Eng. degree from Korea University in 2009, and the Ph.D. degree from Purdue University in 2015, both in electrical & computer engineering. During his Ph.D., he worked with Intel Labs, Samsung Advanced Institute of Technology, and Neuroscience Research Institute, as a Research Intern or a Visiting Lecturer. Prior to joining SKKU, he was a Post-doctoral Research Associate in Mathematics, Purdue

University, a Research Fellow in EECS, the University of Michigan, and an Assistant Professor of ECE at the University of Hawai'i, Mānoa, from 2015 to 2016, from 2016 to 2019, from 2019 to 2021, respectively. His research interests include artificial intelligence & machine learning, optimization, and compressed sensing, applied to applications in computational imaging, image processing, and computer vision.



Jinkyu Lee was born in Chungju, Korea in 1993. He received B.S. degree in Computer Science from Handong Global University, Pohang, Korea, in 2019. He is currently pursuing M.S. degree in the Department of Information and Communication Engineering at the Handong Global University.

From 2016 to now, he is a research assistant with the computer graphics and vision laboratory. His current research interests are computer vision, simultaneous localization and mapping system, advanced driver assistance system, and autonomous driving.



Muhyun Back was born in Gwangju, Korea in 1994. He received B.S. degree in Computer Science from Handong Global University, Pohang, Korea, in 2019. He is currently pursuing M.S. degree in the Department of Information and Communication Engineering at the Handong Global University.

From 2016 to now, he is a research assistant with the computer graphics and vision laboratory. His current research interests are computer vision, advanced driver assistance system, simultaneous localization and mapping system, and 3D reconstruction.