VDO-SLAM: A Visual Dynamic Object-aware SLAM System

Jun Zhang[co]1, Mina Henein[co]1, Robert Mahony1 and Viorela Ila2

Abstract

The scene rigidity assumption, also known as the static world assumption, is common in SLAM algorithms. Most existing algorithms operating in complex dynamic environments simplify the problem by removing moving objects from consideration or tracking them separately. Such strong assumptions limit the deployment of autonomous mobile robotic systems in a wide range of important real world applications involving highly dynamic and unstructured environments. This paper presents VDO-SLAM, a robust object-aware dynamic SLAM system that exploits semantic information to enable motion estimation of rigid objects in the scene without any prior knowledge of the objects shape or motion models. The proposed approach integrates dynamic and static structures in the environment into a unified estimation framework resulting in accurate robot pose and spatio-temporal map estimation. We provide a way to extract velocity estimates from object pose change of moving objects in the scene providing an important functionality for navigation in complex dynamic environments. We demonstrate the performance of the proposed system on a number of real indoor and outdoor datasets. Results show consistent and substantial improvements over state-of-the-art algorithms. An open-source version of the source code is available*.

Keywords

SLAM, dynamic environment, object-aware, velocity, tracking, pose change

1 Introduction

The demand for high-quality three dimensional virtual models of complex environments is growing in a wide range of applications from robotics, intelligent transportation, surveillance, inspection, entertainment and film production to exploration and monitoring of natural environments, among many others. Creating these models involves repeatedly sensing the environment and fusing sensor measurements into a consistent representation using estimation techniques such as Simultaneous Localisation And Mapping (SLAM) or Structure from Motion (SfM). SLAM, a well established research field in robotics, comprises simultaneously estimating a robot's state and building a model representation of its environment. While many accurate and efficient solutions to the problem exist, current SLAM algorithms can be easily induced to fail in highly dynamic environments (Cadena et al. (2016)). The conventional technique to deal with dynamics in SLAM is to either treat any sensor data associated with moving objects as outliers and remove them from the estimation process (Hahnel et al. (2002, 2003); Wolf and Sukhatme (2005); Zhao et al. (2008); Bescos et al. (2018)), or detect moving objects and track them separately using traditional multi-target tracking approaches (Wang et al. (2003); Miller and Campbell (2007); Rogers et al. (2010); Kundu et al. (2011)). The former technique excludes information about dynamic objects in the scene, and generates static only maps. Accuracy of the latter is dependent on the camera pose estimation, which is more susceptible to failure in complex dynamic environments where the presence of reliable static structure is questionable. Increased applications of autonomous systems to dynamic

environments is driving the community to challenge the scene rigidity assumption, also known as the static world assumption, that underpins most existing open-source SLAM and Visual Odometry (VO) algorithms.

A typical SLAM system consists of a front-end module, that processes the raw data from the sensors and a back-end module, that integrates the obtained information (raw and implicit higher-level information) into a probabilistic estimation framework. Simple primitives such as 3D locations of salient features are commonly used to represent the environment. This is largely a consequence of the fact that points are easy to detect, track and integrate within the SLAM estimation problem. Other primitives such as lines and planes (de la Puente and Rodríguez-Losada (2014); Kaess (2015); Henein et al. (2017); Hsiao et al. (2017)) or even objects (Mu et al. (2016); Salas-Moreno et al. (2013); Yang and Scherer (2019)) have been considered in order to provide richer map representations. Semantic information and object segmentation can provide important prior information in identifying dynamic objects in the scene (Wang et al. (2007); Gálvez-López et al. (2016)).

[co] The two authors contributed equally to this work.

* https://github.com/halajun/vdo_slam

1 Australian National University, Canberra
2 University of Sydney, Sydney

Corresponding author:
Jun Zhang, Systems Theory and Robotics Laboratory, 115 North Road, 2601 Acton, ACT, Australia.
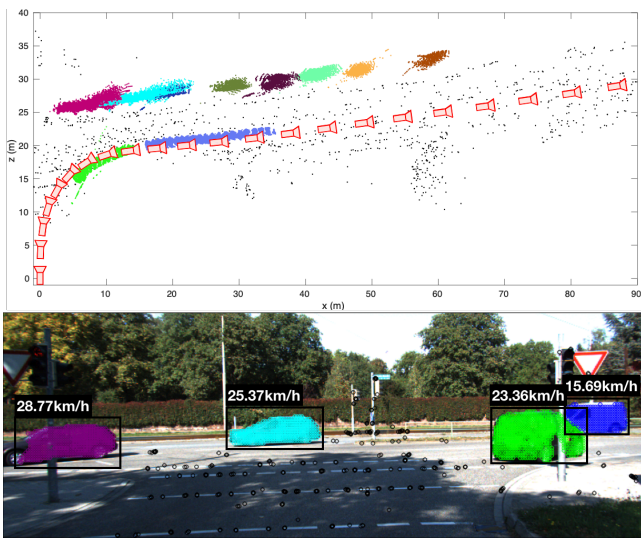Email: jun.zhang2@anu.edu.au

Figure 1. Results of our VDO-SLAM system. (Top) A full map including camera trajectory, static background and moving object structure. (Bottom) Detected points on static background and object body, and estimated object speed. Black circles represents static points, and each object is shown with a different colour.

Advances in deep learning have provided algorithms that can reliably detect and segment classes of objects at almost real time (Girshick et al. (2018); He et al. (2017)). Despite recent developments in vision-based object detection and segmentation, the visual SLAM community has not yet fully exploited such information (Nicholson et al. (2018)). To incorporate such information in existing geometric SLAM algorithms then either a dataset of 3D-models of every object in the scene must be available (Salas-Moreno et al. (2013); Gálvez-López et al. (2016)) or the front end must explicitly provide object pose information in addition to detection and segmentation (Milan et al. (2016); Byravan and Fox (2017); Wohlhart and Lepetit (2015)) adding a layer of complexity to the problem. The requirement for accurate 3D-models severely limits the potential domains of application, while to the best of our knowledge, multiple object tracking and 3D pose estimation remain a challenge to learning techniques. There is a clear need for an algorithm that can exploit the powerful detection and segmentation capabilities of modern deep learning algorithms without relying on additional pose estimation or object model prior.

In this paper, we propose VDO-SLAM, a novel feature-based stereo/RGB-D dynamic SLAM system, that leverages image-based semantic information to simultaneously localise the robot, map the static and dynamic structure, and track motions of rigid objects in the scene. In summary, the contributions of this work are:

- a novel formulation to model dynamic scenes in a unified estimation framework over robot poses, static and dynamic 3D points, and object motions.

- accurate estimation for SE(3) pose change of dynamic objects that outperforms state-of-the-art algorithms, as well as a way to extract objects' velocity in the scene,

- a robust method for tracking moving objects exploiting semantic information with the ability to handle

- indirect occlusions resulting from the failure of semantic object segmentation,

- a demonstrable system in complex and compelling real-world scenarios.

To the best of our knowledge, this is the first full dynamic SLAM system that is able to achieve motion segmentation, dynamic object tracking, and estimate the camera poses along with the static and dynamic structure, the full SE(3) pose change of every rigid object in the scene, extract velocity information, and be demonstrable in real-world outdoor scenarios. We demonstrate the performance of our algorithm on real datasets and show capability of the proposed system to resolve rigid object motion estimation and yield motion results that are comparable to the camera pose estimation in accuracy and that outperform state-of-the-art algorithms by an order of magnitude in urban driving scenarios.

The remainder of this paper is structured as follows, in the following Section 2 we discuss the related work. In Section 3 and 4 we describe the proposed algorithm and system. We introduce the experimental setup, followed by the results and evaluations in Section 5. We summarise and offer concluding remarks in Section 6.

2 Related Work

In order to perform tasks autonomously, a robot must be able to reason about its environment. Establishing the spatial and temporal relationships between a robot, stationary and moving objects in a scene serves as a basis for scene understanding and the problems of simultaneous localisation, mapping and moving object tracking are mutually beneficial (Wang et al. (2007)). In the SLAM community, however, information associated with stationary objects is considered positive, while information drawn from moving objects is seen as degrading the algorithm performance. State-of-the-art SLAM systems either treat data from moving objects as outliers (Hahnel et al. (2002, 2003); Wolf and Sukhatme (2005); Zhao et al. (2008); Bescos et al. (2018)) or track them separately using multi-target tracking (Wang et al. (2003); Miller and Campbell (2007); Rogers et al. (2010); Kundu et al. (2011)), and very few aim to utilise information from static and dynamic objects into a single framework to improve the accuracy of the estimation (Bibby and Reid (2007); Judd et al. (2018); Yang and Scherer (2019)).

One of the earliest works in the area of SLAM in dynamic environments is presented by Hahnel et al. (2003) who use an Expectation-Maximisation (EM) algorithm to update the probabilistic estimate about which measurements correspond to a static/dynamic object, and remove them from the estimation when they correspond to a dynamic object. Alcantarilla et al. (2012) introduce dense scene flow for dynamic objects detection, and show improved localisation and mapping results by removing "erroneous" measurements on dynamic objects from the estimation. Tan et al. (2013) propose an online key-frame update that reliably detects changed features in terms of appearance and structure and discards them if necessary.

Wang et al. (2007) developed a theory for performing SLAM with Moving Objects Tracking (SLAMMOT). In the

latest version of their SLAM with detection and tracking of moving objects, the estimation problem is decomposed into two separate estimators (moving and stationary objects) to make it feasible to update both filters in real time. Kundu *et al.* (Kundu et al. (2011)) tackle the SLAM problem with dynamic objects by solving the problems of Structure from Motion (SfM) and tracking of moving objects in parallel, and unifying the output of the system into a 3D dynamic map of the scene containing the structure and the trajectory of both static and moving objects. Reddy et al. (2015) use optical flow and depth to compute semantic motion segmentation. They isolate static objects from moving objects and reconstruct them independently, before using semantic constraints to improve the 3D reconstruction.

Bibby and Reid's SLAMIDE (Bibby and Reid (2007)) estimates the state of 3D features (stationary or dynamic) with a generalised EM algorithm where they use reversible data association to include dynamic objects in a single framework SLAM. A multi-camera SLAM system is proposed by Zou and Tan (2013), that is able to track multiple cameras, as well as reconstruct the 3D position of both static background and moving foreground points. Their system leverages the idea that points on moving objects share information about relative camera poses at the same time step to estimate all camera poses simultaneously. Judd et al. (2018) estimate the full SE(3) motion of both the camera and rigid objects in the scene by applying a multi-motion visual odometry (MVO) multi-model fitting technique. A very recent work by Yang and Scherer (2019) presents a method for single image 3D cuboid detection, and multi-view object SLAM for both static and dynamic environments. Their main interest, however, is the camera pose and object detection accuracy and they provide no evaluation of the accuracy of their object pose estimation.

In the last two decades, significant research has been made in the computer vision community to solve the problems of object motion segmentation and multiple object tracking, which are important components in estimating dynamic scenes. Current state-of-the-art motion segmentation algorithms seek to combine methods based on affine assumption (Vidal and Hartley (2004)) and those based on Epipolar geometry (Rao et al. (2010)) into a single framework that leverages their advantages (Elhamifar and Vidal (2009); Lai et al. (2016)). In a recent approach, a multi-frame spectral clustering framework is introduced by Xu et al. (2018) with joint integration of an affine model, a homography model and a fundamental matrix. Making use of semantic information has proven to help deal with the issues of degenerate motions and partial occlusions increasing the motion segmentation accuracy (Rubino et al. (2018)). Multiple object tracking algorithms have moved over the last few years from classical inference/filtering based (Khan et al. (2005); Zhou et al. (2015)) to data-driven (deep learning) approaches (Milan et al. (2017); Kim et al. (2018)). The state-of-the-art STAM-MOT (Chu et al. (2017)) applies spatial and temporal attention map to handle the partial occlusion problem in tracking. To find the optimal location of objects, the proposed algorithm employs the dense searching strategies, which are utilised commonly in tracking single object.

3 Methodology

In this section we show how to model the motion of a rigid-object in a model free manner based on point tracking. We propose a factor graph optimisation to estimate the camera and object motion.

In the tracking component of our system, shown in Fig. 4, the cost function chosen to estimate the camera pose and object motion (described in Section 3.2) is associated with the 3D-2D re-projection error and is defined on the image plane. Since the noise is better characterised in image plane, this yields more accurate results for camera localisation (Nistér et al. (2004)). Moreover, based on this error term, we propose a novel formulation to jointly optimise the optical flow along with the camera pose and the object motion, to ensure a robust tracking of points (described in Section 3.2.3). In the mapping module, a 3D error cost function is used to ensure best results of 3D structure and object motions estimation as described in Section 3.3.

3.1 Background and Notation

3.1.1 Coordinate Frames: Let $^0\mathbf{X}_k, ^0\mathbf{L}_k \in \text{SE}(3)$ be the robot/camera and the object 3D pose respectively, at time k in a global reference frame 0, with $k \in \mathscr{T}$ the set of time steps. Note that calligraphic capital letters are used in our notation to represent sets of indices. Fig. 2 shows these pose transformations as solid curves.

3.1.2 Points: Let $^0\mathbf{m}_k^i$ be the homogeneous coordinates of the i^{th} 3D point at time k, with $^0\mathbf{m}^i = \left[m_x^i, m_y^i, m_z^i, 1\right]^\top \in \mathbb{IE}^3$ and $i \in \mathscr{M}$ the set of points. We write a point in robot/camera frame as $^{X_k}\mathbf{m}_k^i = ^0\mathbf{X}_k^{-1}\,^0\mathbf{m}_k^i$.

Define $\mathbf{I}_k$ the reference frame associated with the image captured by the camera at time k chosen at the top left corner of the image, and let $^{I_k}\mathbf{p}_k^i = \left[u^i, v^i, 1\right] \in \mathbb{IE}^2$ be the pixel location on frame $\mathbf{I}_k$ corresponding to the homogeneous 3D point $^{X_k}\mathbf{m}_k^i$, which is obtained via the projection function $\pi(\cdot)$ as follows:

$$^{I_k}\mathbf{p}_k^i = \pi(^{X_k}\mathbf{m}_k^i) = \mathbf{K}\,^{X_k}\mathbf{m}_k^i, \tag{1}$$

where $\mathbf{K}$ is the camera intrinsics matrix.

The camera and/or object motions both produce an optical flow $^{I_k}\boldsymbol{\phi}^i \in \mathbb{R}^2$ that is the displacement vector indicating the motion of pixel $^{I_{k-1}}\mathbf{p}_{k-1}^i$ from image frame I_{k-1} to I_k, and is given by:

$$^{I_k}\boldsymbol{\phi}^i = ^{I_k}\tilde{\mathbf{p}}_k^i - ^{I_{k-1}}\mathbf{p}_{k-1}^i. \tag{2}$$

Here $^{I_k}\tilde{\mathbf{p}}_k^i$ is the correspondence of $^{I_{k-1}}\mathbf{p}_{k-1}^i$ in I_k. Note that, we overload the same notation to represent the 2D pixel coordinates $\in \mathbb{R}^2$. In this work, we leverage optical flow to find correspondences between consecutive frames.

3.1.3 Object and 3D Point Motions: The object motion between times $k-1$ and k is described by the homogeneous transformation $^{L_{k-1}}_{k-1}\mathbf{H}_k \in \text{SE}(3)$ according to:

$$^{L_{k-1}}_{k-1}\mathbf{H}_k = ^0\mathbf{L}_{k-1}^{-1}\,^0\mathbf{L}_k. \tag{3}$$

Fig. 2 shows these motion transformations as dashed curves. We write a point in its corresponding object frame as $^{L_k}\mathbf{m}_k^i = ^0\mathbf{L}_k^{-1}\,^0\mathbf{m}_k^i$ (shown as a dashed vector from the object
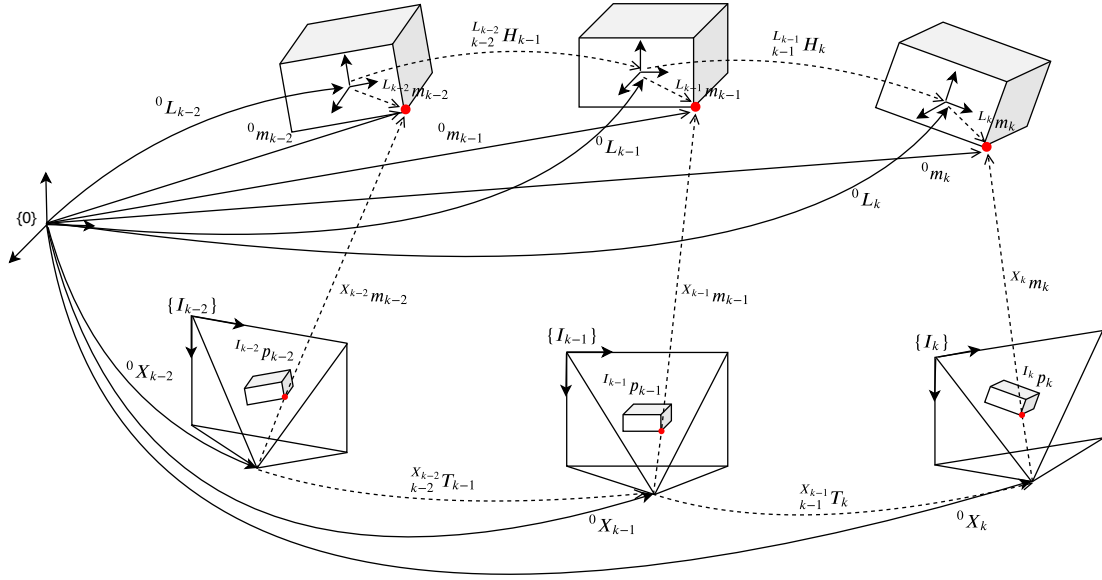
Figure 2. Notation and coordinate frames. Solid curves represent camera and object poses in inertial frame; $^0\mathbf{X}$ and $^0\mathbf{L}$ respectively, and dashed curves their respective motions in body-fixed frame. Solid lines represent 3D points in inertial frame, and dashed lines represent 3D points in camera frames.

reference frame to the red dot in Fig. 2), substituting the object pose at time k from (3), this becomes:

$$^0\mathbf{m}_k^i = {}^0\mathbf{L}_k\, {}^{L_k}\mathbf{m}_k^i = {}^0\mathbf{L}_{k-1}\, {}^{L_{k-1}}_{k-1}\mathbf{H}_k\, {}^{L_k}\mathbf{m}_k^i\,. \qquad (4)$$

Note that for rigid body objects, $^{L_k}\mathbf{m}_k^i$ stays constant at $^L\mathbf{m}^i$, and $^L\mathbf{m}^i = {}^0\mathbf{L}_k^{-1}\, {}^0\mathbf{m}_k^i = {}^0\mathbf{L}_{k+n}^{-1}\, {}^0\mathbf{m}_{k+n}^i$ for any integer $n \in \mathbb{Z}$. Then, for rigid objects with $n = -1$, (4) becomes:

$$^0\mathbf{m}_k^i = {}^0\mathbf{L}_{k-1}\, {}^{L_{k-1}}_{k-1}\mathbf{H}_k\, {}^0\mathbf{L}_{k-1}^{-1}\, {}^0\mathbf{m}_{k-1}^i\,. \qquad (5)$$

(5) is crucially important as it relates the same 3D point on a rigid object in motion at consecutive time steps by a homogeneous transformation $^0_{k-1}\mathbf{H}_k := {}^0\mathbf{L}_{k-1}\, {}^{L_{k-1}}_{k-1}\mathbf{H}_k\, {}^0\mathbf{L}_{k-1}^{-1}$. This equation represents a *frame change of a pose transformation* (Chirikjian et al. (2017)), and shows how the body-fixed frame pose change $^{L_{k-1}}_{k-1}\mathbf{H}_k$ relates to the global reference frame pose change $^0_{k-1}\mathbf{H}_k$. The point motion in global reference frame is then expressed as:

$$^0\mathbf{m}_k^i = {}^0_{k-1}\mathbf{H}_k\, {}^0\mathbf{m}_{k-1}^i\,. \qquad (6)$$

Equation (6) is at the core of our motion estimation approach, as it expresses the rigid object pose change in terms of the points that reside on the object in a model-free manner without the need to include the object 3D pose as a random variable in the estimation. Section 3.2.2 details how this rigid object pose change is estimated based on the above equation. Here $^0_{k-1}\mathbf{H}_k \in SE(3)$ represents the object point motion in global reference frame; for the remainder of this document, we refer to this quantity as the object pose change or the object motion for ease of reading.

3.2 Camera Pose and Object Motion Estimation

3.2.1 Camera Pose Estimation:
Given a set of static 3D points $\{^0\mathbf{m}_{k-1}^i \mid i \in \mathcal{M}, k \in \mathcal{T}\}$ observed at time $k-1$ in

global reference frame, and the set of 2D correspondences $\{^{I_k}\tilde{\mathbf{p}}_k^i \mid i \in \mathcal{M}, k \in \mathcal{T}\}$ in image $\mathbf{I}_k$, the camera pose $^0\mathbf{X}_k$ is estimated via minimizing the re-projection error:

$$\mathbf{e}_i(^0\mathbf{X}_k) = {}^{I_k}\tilde{\mathbf{p}}_k^i - \pi(^0\mathbf{X}_k^{-1}\, {}^0\mathbf{m}_{k-1}^i)\,. \qquad (7)$$

We parameterise the $SE(3)$ camera pose by elements of the Lie-algebra $\mathbf{x}_k \in se(3)$:

$$^0\mathbf{X}_k = \exp(^0\mathbf{x}_k)\,, \qquad (8)$$

and define $^0\mathbf{x}_k^\vee \in \mathbb{R}^6$ with the *vee* operator a mapping from $se(3)$ to $\mathbb{R}^6$. Using the Lie-algebra parameterisation of $SE(3)$ with the substitution of (8) into (7), the solution of the least squares cost is given by:

$$^0\mathbf{x}_k^{*\vee} = \underset{^0\mathbf{x}_k^\vee}{\operatorname{argmin}} \sum_i^{n_b} \rho_h\left(\mathbf{e}_i^\top(^0\mathbf{x}_k)\,\Sigma_p^{-1}\,\mathbf{e}_i(^0\mathbf{x}_k)\right) \qquad (9)$$

for all n_b visible 3D-2D static background point correspondences between consecutive frames. Here ρ_h is the Huber function (Huber (1992)), and Σ_p is the covariance matrix associated with the re-projection error. The estimated camera pose is given by $^0\mathbf{X}_k^* = \exp(^0\mathbf{x}_k^*)$ and is found using the Levenberg-Marquardt algorithm to solve for (9).

3.2.2 Object Motion Estimation:
Analogous to the camera pose estimation, a cost function based on re-projection error is constructed to solve for the object motion $^0_{k-1}\mathbf{H}_k$. Using (6), the error term between the re-projection of an object 3D point and the corresponding 2D point in image $\mathbf{I}_k$ is:

$$\begin{aligned}\mathbf{e}_i(^0_{k-1}\mathbf{H}_k) :=&\ {}^{I_k}\tilde{\mathbf{p}}_k^i - \pi(^0\mathbf{X}_k^{-1}\, {}^0_{k-1}\mathbf{H}_k\, {}^0\mathbf{m}_{k-1}^i)\\ =&\ {}^{I_k}\tilde{\mathbf{p}}_k^i - \pi(^0_{k-1}\mathbf{G}_k\, {}^0\mathbf{m}_{k-1}^i)\,,\end{aligned} \qquad (10)$$

where $^0_{k-1}\mathbf{G}_k \in SE(3)$. Parameterising $^0_{k-1}\mathbf{G}_k := \exp\left(^0_{k-1}\mathbf{g}_k\right)$ with $^0_{k-1}\mathbf{g}_k \in se(3)$, the optimal solution is found via

minimising:

$$_{k-1}^{0}\mathbf{g}_k^{*\vee} = \underset{_{k-1}^{0}\mathbf{g}_k^{\vee}}{\operatorname{argmin}} \sum_i^{n_d} \rho_h\left(\mathbf{e}_i^{\top}(_{k-1}^{0}\mathbf{g}_k)\,\Sigma_p^{-1}\,\mathbf{e}_i(_{k-1}^{0}\mathbf{g}_k)\right) \quad (11)$$

given all n_d visible 3D-2D dynamic point correspondences on an object between frames $k-1$ and k. The object motion, $_{k-1}^{0}\mathbf{H}_k = {}^0\mathbf{X}_k\,_{k-1}^{0}\mathbf{G}_k$ can be recovered afterwards.

3.2.3 Joint Estimation with Optical Flow: The camera pose and object motion estimation both rely on good image correspondences. Tracking of points on moving objects can be very challenging due to occlusions, large relative motions and large camera-object distances. In order to ensure a robust tracking of points, the technique proposed in this paper aims at refining the estimation of the optical flow jointly with the motion estimation.

For camera pose estimation, the error term in (7) is reformulated considering (2) as:

$$\mathbf{e}_i({}^0\mathbf{X}_k, {}^{I_k}\boldsymbol{\phi}) = {}^{I_{k-1}}\mathbf{p}_{k-1}^i + {}^{I_k}\boldsymbol{\phi}^i - \pi({}^0\mathbf{X}_k^{-1}\,{}^0\mathbf{m}_{k-1}^i)\,. \quad (12)$$

Applying the Lie-algebra parameterisation of SE(3) element, the optimal solution is obtained via minimising the cost function:

$$\{{}^0\mathbf{x}_k^{*\vee}, {}^k\boldsymbol{\Phi}_k^*\} = \underset{\{{}^0\mathbf{x}_k^{\vee}, {}^k\boldsymbol{\Phi}_k\}}{\operatorname{argmin}} \sum_i^{n_b} \Big\{ \rho_h\big(\mathbf{e}_i^{\top}({}^{I_k}\boldsymbol{\phi}^i)\,\Sigma_\phi^{-1}\,\mathbf{e}_i({}^{I_k}\boldsymbol{\phi}^i)\big) +$$
$$\rho_h\big(\mathbf{e}_i^{\top}({}^0\mathbf{x}_k, {}^{I_k}\boldsymbol{\phi}^i)\,\Sigma_p^{-1}\,\mathbf{e}_i({}^0\mathbf{x}_k, {}^{I_k}\boldsymbol{\phi}^i)\big)\Big\}\,, \quad (13)$$

where $\rho_h(\mathbf{e}_i^{\top}({}^{I_k}\boldsymbol{\phi}^i)\,\Sigma_\phi^{-1}\,\mathbf{e}_i({}^{I_k}\boldsymbol{\phi}^i))$ is the regularization term with

$$\mathbf{e}_i({}^{I_k}\boldsymbol{\phi}^i) = {}^{I_k}\hat{\boldsymbol{\phi}}^i - {}^{I_k}\boldsymbol{\phi}^i\,. \quad (14)$$

Here ${}^{I_k}\hat{\boldsymbol{\Phi}}^i = \{{}^{I_k}\hat{\boldsymbol{\phi}}^i \mid i \in \mathscr{M}, k \in \mathscr{T}\}$ is the initial optic-flow obtained through classical or learning-based methods, and Σ_ϕ is the associated covariance matrix. Analogously, the cost function for object motion in (11) combining optical flow refinement is given by

$$\{{}_{k-1}^{0}\mathbf{g}_k^{*\vee}, {}^k\boldsymbol{\Phi}_k^*\} = \underset{\{{}_{k-1}^{0}\mathbf{g}_k^{\vee}, {}^k\boldsymbol{\Phi}_k\}}{\operatorname{argmin}} \sum_i^{n_d} \Big\{ \rho_h\big(\mathbf{e}_i^{\top}({}^{I_k}\boldsymbol{\phi}^i)\,\Sigma_\phi^{-1}\,\mathbf{e}_i({}^{I_k}\boldsymbol{\phi}^i)\big) +$$
$$\rho_h\big(\mathbf{e}_i^{\top}(_{k-1}^{0}\mathbf{g}_k, {}^{I_k}\boldsymbol{\phi}^i)\,\Sigma_p^{-1}\,\mathbf{e}_i(_{k-1}^{0}\mathbf{g}_k, {}^{I_k}\boldsymbol{\phi}^i)\big)\Big\}\,. \quad (15)$$

3.3 Graph Optimisation

The proposed approach formulates the dynamic SLAM as a graph optimisation problem, to refine the camera poses and object motions, and build a global consistent map including static and dynamic structure. We model the dynamic SLAM problem as a factor graph as the one demonstrated in Fig. 3. The factor graph formulation is highly intuitive and has the advantage that it allows for efficient implementations of batch (Dellaert and Kaess (2006); Agarwal et al. (2012)) and incremental (Kaess et al. (2011); Polok et al. (2013); Ila et al. (2017)) solvers.

Four types of measurements/observations are integrated into a joint optimisation problem; the 3D point measurements, the visual odometry measurements, the motion

of points on a dynamic object and the object smooth motion observations. The 3D point measurement model error $\mathbf{e}_{i,k}({}^0\mathbf{X}_k, {}^0\mathbf{m}_k^i)$ is defined as:

$$\mathbf{e}_{i,k}({}^0\mathbf{X}_k, {}^0\mathbf{m}_k^i) = {}^0\mathbf{X}_k^{-1}\,{}^0\mathbf{m}_k^i - \mathbf{z}_k^i\,. \quad (16)$$

Here $\mathbf{z} = \{\mathbf{z}_k^i \mid i \in \mathscr{M}, k \in \mathscr{T}\}$ is the set of all 3D point measurements at all time steps, with cardinality n_z and $\mathbf{z}_k^i \in \mathbb{R}^3$. The 3D point measurement factors are shown as white circles in Fig. 3.

The visual odometry model error $\mathbf{e}_k({}^0\mathbf{X}_{k-1}, {}^0\mathbf{X}_k)$ is defined as:

$$\mathbf{e}_k({}^0\mathbf{X}_{k-1}, {}^0\mathbf{X}_k) = ({}^0\mathbf{X}_{k-1}^{-1}\,{}^0\mathbf{X}_k)^{-1}\,_{k-1}^{X_{k-1}}\mathbf{T}_k\,, \quad (17)$$

where $\mathbf{T} = \{{}_{k-1}^{X_{k-1}}\mathbf{T}_k \mid k \in \mathscr{T}\}$ is the odometry measurement set with ${}_{k-1}^{X_{k-1}}\mathbf{T}_k \in$ SE(3) and cardinality n_o. The odometric factors are shown as orange circles in Fig. 3.

The motion model error of points on dynamic objects $\mathbf{e}_{i,l,k}({}^0\mathbf{m}_k^i, {}_{k-1}^{0}\mathbf{H}_k^l, {}^0\mathbf{m}_{k-1}^i)$ is defined as:

$$\mathbf{e}_{i,l,k}({}^0\mathbf{m}_k^i, {}_{k-1}^{0}\mathbf{H}_k^l, {}^0\mathbf{m}_{k-1}^i) = {}^0\mathbf{m}_k^i - {}_{k-1}^{0}\mathbf{H}_k^l\,{}^0\mathbf{m}_{k-1}^i\,. \quad (18)$$

The motion of all points on a detected rigid object l are characterised by the same pose transformation ${}_{k-1}^{0}\mathbf{H}_k^l \in$ SE(3) given by (6) and the corresponding factor, shown as magenta circles in Fig. 3, is a ternary factor which we call the *motion model of a point on a rigid body*.

It has been shown that incorporating prior knowledge about the motion of objects in the scene is highly valuable in dynamic SLAM (Wang et al. (2007); Henein et al. (2020)). Motivated by the camera frame rate and the physics laws governing the motion of relatively large objects (vehicles) and preventing their motions to change abruptly, we introduce smooth motion factors to minimise the change in consecutive object motions, with the error term defined as:

$$\mathbf{e}_{l,k}({}_{k-2}^{0}\mathbf{H}_{k-1}^l, {}_{k-1}^{0}\mathbf{H}_k^l) = {}_{k-2}^{0}\mathbf{H}_{k-1}^l{}^{-1}\,{}_{k-1}^{0}\mathbf{H}_k^l\,. \quad (19)$$

The object smooth motion factor $\mathbf{e}_{l,k}({}_{k-2}^{0}\mathbf{H}_{k-1}^l, {}_{k-1}^{0}\mathbf{H}_k^l)$ is used to minimise the change between the object motion at consecutive time steps and is shown as cyan circles in Fig. 3.

Let $\boldsymbol{\theta}_M = \{{}^0\mathbf{m}_k^i \mid i \in \mathscr{M}, k \in \mathscr{T}\}$ be the set of all 3D points, and $\boldsymbol{\theta}_X = \{{}^0\mathbf{x}_k^{\vee} \mid k \in \mathscr{T}\}$ as the set of all camera poses. We parameterise the SE(3) object motion ${}_{k-1}^{0}\mathbf{H}_k^l$ by elements ${}_{k-1}^{0}\mathbf{h}_k^l \in$ se(3) the Lie-algebra of SE(3):

$$_{k-1}^{0}\mathbf{H}_k^l = \exp(_{k-1}^{0}\mathbf{h}_k^l)\,, \quad (20)$$

and define $\boldsymbol{\theta}_H = \{{}_{k-1}^{0}\mathbf{h}_k^{l\,\vee} \mid k \in \mathscr{T}, l \in \mathscr{L}\}$ as the set of all object motions, with ${}_{k-1}^{0}\mathbf{h}_k^{l\,\vee} \in \mathbb{R}^6$ and $\mathscr{L}$ the set of all object labels. Given $\boldsymbol{\theta} = \boldsymbol{\theta}_X \cup \boldsymbol{\theta}_M \cup \boldsymbol{\theta}_H$ as all the nodes in the graph, and using the Lie-algebra parameterisation of SE(3) for $\mathbf{X}$ and $\mathbf{H}$ (substituting (8) in (16) and (17), and substituting (20) in (18) and (19)), the solution of the least

squares cost is given by:

$$
\begin{aligned}
\boldsymbol{\theta}^* = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \Big\{ & \sum_{i,k}^{n_z} \rho_h \big(\mathbf{e}_{i,k}^\top ({}^0\mathbf{x}_k, {}^0\mathbf{m}_k^i) \, \Sigma_z^{-1} \, \mathbf{e}_{i,k} ({}^0\mathbf{x}_k, {}^0\mathbf{m}_k^i) \big) \\
+ & \sum_k^{n_o} \rho_h \big(\log(\mathbf{e}_k ({}^0\mathbf{x}_{k-1}, {}^0\mathbf{x}_k))^\top \Sigma_o^{-1} \, \log(\mathbf{e}_k ({}^0\mathbf{x}_{k-1}, {}^0\mathbf{x}_k)) \big) \\
+ & \sum_{i,l,k}^{n_g} \rho_h \big(\mathbf{e}_{i,l,k}^\top ({}^0\mathbf{m}_k^i, {}_{k-1}{}^0\mathbf{h}_k^l, {}^0\mathbf{m}_{k-1}^i) \, \Sigma_g^{-1} \\
& \qquad \mathbf{e}_{i,l,k} ({}^0\mathbf{m}_k^i, {}_{k-1}{}^0\mathbf{h}_k^l, {}^0\mathbf{m}_{k-1}^i) \big) \\
+ & \sum_{l,k}^{n_s} \rho_h \big(\log(\mathbf{e}_{l,k} ({}_{k-2}{}^0\mathbf{h}_{k-1}^l, {}_{k-1}{}^0\mathbf{h}_k^l))^\top \Sigma_s^{-1} \\
& \qquad\qquad \log(\mathbf{e}_{l,k} ({}_{k-2}{}^0\mathbf{h}_{k-1}^l, {}_{k-1}{}^0\mathbf{h}_k^l))) \Big\}, \quad (21)
\end{aligned}
$$

where Σ_z is the 3D point measurement noise covariance matrix, Σ_o is the odometry noise covariance matrix, Σ_g is the motion noise covariance matrix with n_g the total number of ternary object motion factors, and Σ_s the smooth motion covariance matrix, with n_s the total number of smooth motion factors. The non-linear least squares problem in (21) is solved using Levenberg-Marquardt method.
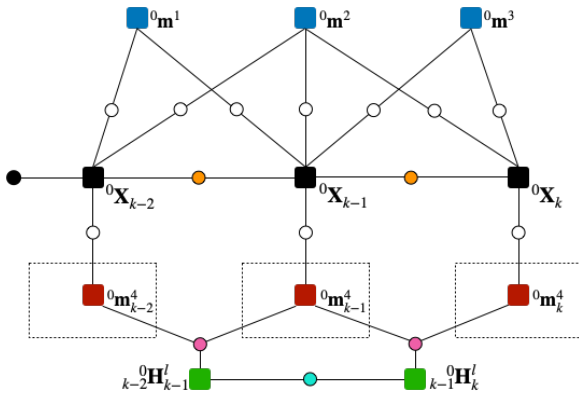

Figure 3. Factor graph representation of an object-aware SLAM with a moving object. Black squares represent the camera poses at different time steps, blue squares represent three static points, red squares represent the same dynamic point on an object (dashed box) at different time steps and green squares the object pose change between time steps. For ease of visualisation, only one dynamic point is drawn here, however, at the time of estimation, all points on a detected dynamic object are used. A prior unary factor is shown as a black circle, odometry binary factors are shown as orange circles, point measurement binary factors as white circles and point motion ternary factors as magenta circles. A smooth motion binary factor is shown as cyan circle.

4 System

In this section, we propose a novel object-aware dynamic SLAM system that robustly estimates both camera and object motions, along with the static and dynamic structure of the environment. The full system overview is shown in Fig. 4. The system consists of three main components: image pre-processing, tracking and mapping.

The input to the system is stereo or RGB-D images. For stereo images, as a first step, we extract depth information by applying the stereo depth estimation method described

in Yamaguchi et al. (2014) to generate depth maps and the resulting data is treated as RGB-D.

Although this system was initially designed to be an RGB-D system, as an attempt to fully exploit image-based semantic information, we apply single image depth estimation to achieve depth information from monocular camera. Our "learning-based monocular" system is monocular in the sense that only RGB images are used as input to the system, however the estimation problem is formulated using RGB-D data, where the depth is obtained using single image depth estimation.

4.1 Pre-processing

There are two challenging aspects that this module needs to fulfil. First, to robustly separate static background and objects, and secondly to ensure long-term tracking of dynamic objects. To achieve this, we leverage recent advances in computer vision techniques for instance level semantic segmentation and dense optical flow estimation in order to ensure efficient object motion segmentation and robust object tracking.

4.1.1 Object Instance Segmentation: Instance-level semantic segmentation is used to segment and identify potentially movable objects in the scene. Semantic information constitutes an important prior in the process of separating static and moving object points, e.g., buildings and roads are always static, but cars can be static or dynamic. Instance segmentation helps to further divide semantic foreground into different instance masks, which makes it easier to track each individual object. Moreover, segmentation masks provide a "precise" boundary of the object body that ensures robust tracking of points on the object.

4.1.2 Optical Flow Estimation: The dense optical flow is used to maximise the number of tracked points on moving objects. Most of the moving objects only occupy a small portion of the image. Therefore, using sparse feature matching does not guarantee robust nor long-term feature tracking. Our approach makes use of dense optical flow to considerably increase the number of object points by sampling from all the points within the semantic mask. Dense optical flow is also used to consistently track multiple objects by propagating a unique object identifier assigned to every point on an object mask. Moreover, it allows to recover objects masks if semantic segmentation fails; a task that is extremely difficult to achieve using sparse feature matching.

4.2 Tracking

The tracking component includes two modules; the camera ego-motion tracking with sub-modules of feature detection and camera pose estimation, and the object motion tracking including sub-modules of dynamic object tracking and object motion estimation.

4.2.1 Feature Detection: To achieve fast camera pose estimation, we detect a *sparse* set of corner features and track them with optical flow. At each frame, only inlier feature points that fit the estimated camera motion are saved into the map, and used to track correspondences in the next frame. New features are detected and added, if the number of inlier
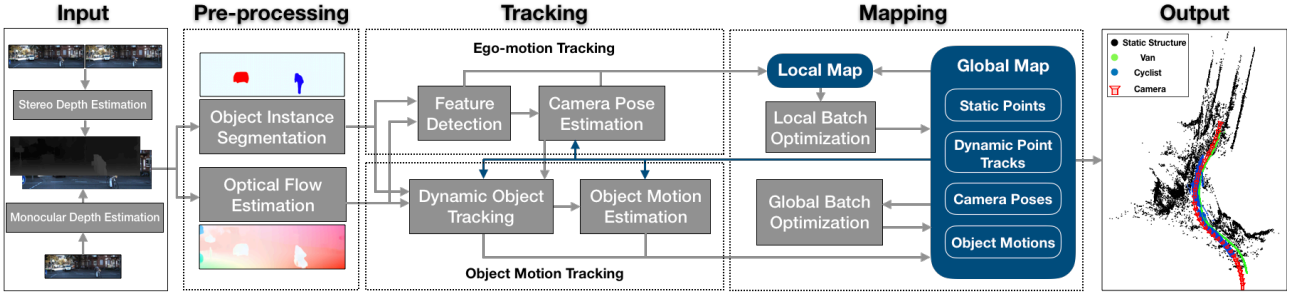
Figure 4. Overview of our VDO-SLAM system. Input images are first pre-processed to generate instance-level object segmentation and dense optical flow. These are then used to track features on static background structure and dynamic objects. Camera poses and object motions estimated from feature tracks are then refined in a global batch optimisation, and a local map is maintained and updated with every new frame. The system outputs camera poses, static structure, tracks of dynamic objects, and estimates of their pose changes over time.

tracks falls below a certain level. These sparse features are detected on static background; image regions excluding the segmented objects.

4.2.2 Camera Pose Estimation: The camera pose is computed using (13) for all detected 3D-2D static point correspondences. To ensure robust estimation, a motion model generation method is applied for initialisation. Specifically, the method generates two models and compares their inlier numbers based on re-projection error. One model is generated by propagating the camera previous motion, while the other by computing a new motion transform using P3P (Ke and Roumeliotis (2017)) algorithm with RANSAC. The motion model that generates most inliers is then selected for initialisation.

4.2.3 Dynamic Object Tracking: The process of object motion tracking consists of two steps. In the first step, segmented objects are classified into static and dynamic. Then we associate the dynamic objects across pairs of consecutive frames.

• Instance-level object segmentation allows us to separate objects from background. Although the algorithm is capable of estimating the motions of all the segmented objects, dynamic object identification helps reduce computational cost of the proposed system. This is done based on scene flow estimation. Specifically, after obtaining the camera pose $^0\mathbf{X}_k$, the scene flow vector $\mathbf{f}_k^i$ describing the motion of a 3D point $^0\mathbf{m}^i$ between frames $k-1$ and k, can be calculated as in Lv et al. (2018):

$$\mathbf{f}_k^i = {}^0\mathbf{m}_{k-1}^i - {}^0\mathbf{m}_k^i = {}^0\mathbf{m}_{k-1}^i - {}^0\mathbf{X}_k\,{}^{X_k}\mathbf{m}_k^i. \qquad (22)$$

Unlike optical flow, scene flow−ideally only caused by scene motion−can directly decide whether some structure is moving or not. Ideally, the magnitude of the scene flow vector should be zero for all static 3D points. However, noise or error in depth and matching complicates the situation in real scenarios. To robustly handle this, we compute the scene flow magnitude of all the sampled points on each object. If the magnitude of the scene flow of a certain point is greater than a predefined threshold, the point is considered dynamic. This threshold was set to 0.12 in all experiments carried in this work. An object is then recognised dynamic if the proportion of "dynamic" points is above a certain level (30% of total number of points), otherwise static. Thresholds to

identify if an object is dynamic were deliberately chosen as mentioned above, to be more conservative as the system is flexible to model a static object as dynamic and estimate a zero motion at every time step, however, the opposite would degrade the system's performance.

• Instance-level object segmentation only provides single-image object labels. Objects then need to be tracked across frames and their motion models propagated over time. We propose to use optical flow to associate point labels across frames. A point label is the same as the unique object identifier on which the point was sampled. We maintain a finite tracking label set $\mathscr{L} \subset \mathbb{N}$, where $l \in \mathscr{L}$ starts from $l = 1$ for the first detected moving object in the scene. The number of elements in $\mathscr{L}$ increases as more moving objects are being detected. Static objects and background are labelled with $l = 0$.

Ideally, for each detected object in frame k, the labels of all its points should be uniquely aligned with the labels of their correspondences in frame $k - 1$. However, in practice this is affected by the noise, image boundaries and occlusions. To overcome this, we assign all the points with the label that appears most in their correspondences. For a dynamic object, if the most frequent label in the previous frame is 0, it means that the object starts to move, appears in the scene at the boundary, or reappears from occlusion. In this case, the object is assigned a new tracking label.

4.2.4 Object Motion Estimation: As mentioned above, objects normally appear in small portions in the scene, which makes it hard to get sufficient sparse features to track and estimate their motions robustly. We sample every third point within an object mask, and track them across frames. Similar to the camera pose estimation, only inlier points are saved into the map and used for tracking in the next frame. When the number of tracked object points decreases below a certain level, new object points are sampled and added. We follow the same method as discussed in Section 4.2.2 to generate an initial object motion model.

4.3 Mapping

In the mapping component, a global map is constructed and maintained. Meanwhile, a local map is extracted from the global map, which is based on the current time step and a window of previous time steps. Both maps are updated via a batch optimisation process.

4.3.1 Local Batch Optimisation: We maintain and update a local map. The goal of the local batch optimisation is to ensure accurate camera pose estimates are provided to the global batch optimisation. The camera pose estimation has a big influence on the accuracy of the object motion estimation and the overall performance of the algorithm. The local map is built using a fixed-size sliding window containing the information of the last n_w frames, where n_w is the window size. Local maps share some common information; this defines the overlap between the different windows. We choose to only locally optimise the camera poses and static structure within the window size, as locally optimising the dynamic structure does not bring any benefit to the optimisation unless a hard constraint (e.g. a constant object motion) is assumed within the window. However, the system is able to incorporate static and dynamic structure in the local mapping if needed. When a local map is constructed, similarly, a factor graph optimisation is performed to refine all the variables within the local map, and then update them back into the global map.

4.3.2 Global Batch Optimisation: The output of the tracking component and the local batch optimisation consists of the camera pose, the object motions and the inlier structure. These are saved in a global map that is constructed with all the previous time steps and is continually updated with every new frame. A factor graph is constructed based on the global map after all input frames have been processed. To effectively explore the temporal constraints, only points that have been tracked for more than 3 instances are added into the factor graph. The graph is formulated as an optimisation problem as described in Section 3.3. The optimisation results serve as the output of the whole system.

4.3.3 From Mapping to Tracking: Maintaining the map provides history information to the estimate of the current state in the tracking module, as shown in Fig. 4 with blue arrows going from the global map to multiple components in the tracking module of the system. Inlier points from the last frame are leveraged to track correspondences in the current frame and estimate camera pose and object motions. The last camera and object motion also serve as possible prior models to initialise the current estimation as described in Section 4.2.2 and 4.2.4. Furthermore, object points help associate semantic masks across frames to ensure robust tracking of objects, by propagating their previously segmented masks in case of "indirect occlusion" resulting from the failure of semantic object segmentation.

5 Experiments

We evaluate our proposed method in terms of camera motion, object motion and velocity, as well as object tracking performance. The evaluation is done on the Oxford Multimotion Dataset (Judd and Gammell (2019)) for indoor, and KITTI Tracking dataset (Geiger et al. (2013)) for outdoor scenarios. We also compare our results to two state-of-the-art methods, MVO (Judd et al. (2018)) and CubeSLAM (Yang and Scherer (2019)), to prove the better performance of VDO-SLAM.

5.1 System Setup

We adopt a learning-based instance-level object segmentation, Mask R-CNN (He et al. (2017)), to generate object segmentation masks. The model of this method is trained on COCO dataset (Lin et al. (2014)), and is directly used in this work without any fine-tuning. For dense optical flow, we leverage a state-of-the-art method; PWC-Net (Sun et al. (2018)). The model is trained on FlyingChairs dataset (Mayer et al. (2016)), and then fine-tuned on Sintel (Butler et al. (2012)) and KITTI training datasets (Geiger et al. (2012)). To generate depth maps for a "monocular" version of our proposed system, we apply a learning-based monocular depth estimation method, MonoDepth2 (Godard et al. (2019)). The model is trained on Depth Eigen split (Eigen et al. (2014)) excluding the tested data in this paper. Feature detection is done using FAST (Rosten and Drummond (2006)) implemented in Rublee et al. (2011).

5.2 Error Metrics

We use a pose change error metric to evaluate the estimated SE(3) motion, i.e., given a ground truth motion transform $\mathbf{T}$ and a corresponding estimated motion $\hat{\mathbf{T}}$, where $\mathbf{T} \in \text{SE}(3)$ could be either a camera relative pose or an object motion. The pose change error is computed as: $E = \hat{\mathbf{T}}^{-1}\,\mathbf{T}$. The translational error E_t is computed as the L_2 norm of the translational component of E. The rotational error E_r is calculated as the angle of rotation in an axis-angle representation of the rotational component of E. For different camera time steps and different objects in a sequence, we compute the root mean squared error (RMSE) for camera poses and object motions, respectively. The object pose change in body-fixed frame is obtained by transforming the pose change $_{k-1}^{0}\mathbf{H}_k$ in the inertial frame into the body frame using the object pose ground-truth

$$_{k-1}^{L_{k-1}}\mathbf{H}_k = {}^0\mathbf{L}_{k-1}^{-1} \; {}_{k-1}^{0}\mathbf{H}_k \; {}^0\mathbf{L}_{k-1}. \tag{23}$$

We also evaluate the object speed error. The linear velocity of a point on the object, expressed in the inertial frame, can be estimated by applying the pose change $_{k-1}^{0}\mathbf{H}_k$ and taking the difference

$$\mathbf{v} \approx {}^0\mathbf{m}_k^i - {}^0\mathbf{m}_{k-1}^i = \left({}_{k-1}^{0}\mathbf{H}_k - \mathbf{I}_4 \right) {}^0\mathbf{m}_{k-1}^i$$
$$= {}_{k-1}^{0}\mathbf{t}_k - (\mathbf{I}_3 - {}_{k-1}^{0}\mathbf{R}_k)\,{}^0\mathbf{m}_{k-1}^i. \tag{24}$$

To get a more reliable measurement, we average over all points on an object at a certain time. Define $\mathbf{c}_{k-1} := \frac{1}{n}\sum \mathbf{m}_{k-1}^i$ for all n points on an object at time $k-1$. Then

$$\mathbf{v} \approx \frac{1}{n}\sum_{i=1}^{n} \left({}_{k-1}^{0}\mathbf{t}_k - (\mathbf{I}_3 - {}_{k-1}^{0}\mathbf{R}_k)\,{}^0\mathbf{m}_{k-1}^i \right)$$
$$= {}_{k-1}^{0}\mathbf{t}_k - (\mathbf{I}_3 - {}_{k-1}^{0}\mathbf{R}_k)\,\mathbf{c}_{k-1}. \tag{25}$$

Then the speed error E_s between the estimated $\hat{\mathbf{v}}$ and the ground truth $\mathbf{v}$ velocities can be calculated as: $E_s = |\hat{\mathbf{v}}| - |\mathbf{v}|$. For different objects tracked over temporal frames, we also compute the RMSE as an error metric.

5.3 Oxford Multimotion Dataset

The recent Oxford Multimotion Dataset (Judd and Gammell (2019)) contains sequences from a moving stereo or RGB-D

camera sensor observing multiple swinging boxes or toy cars in an indoor scenario. Ground truth trajectories of the camera and moving objects are obtained via a Vicon motion capture system. We only choose the swinging boxes sequences for evaluation, since results of real driving scenarios are evaluated on KITTI dataset. Table 1 shows results compared to the state-of-the-art MVO (Judd et al. (2018)). As MVO is a visual odometry system without global refinement, we switch off the batch optimisation module in our system and generate our results for fair comparison. We use the error metrics described in Section 5.2.

Table 1. Comparison versus MVO (Judd et al. (2018)) for camera and object motion estimation accuracy on the sequence of swinging_4_unconstrained sequence in Oxford Multi-motion dataset. Bold numbers indicate the better results.

	Proposed		MVO	
	E_r(deg)	E_t(m)	E_r(deg)	E_t(m)
Camera	**0.4525**	**0.0163**	1.0742	0.0338
Top-left Swinging Box	**1.0175**	**0.0302**	2.9025	0.0685
Top-right Swinging and rotating Box	**1.3567**	0.0229	1.4540	**0.0212**
Bottom-left Swinging Box	**1.6356**	**0.0290**	2.9765	0.0502
Bottom-right Rotating Box	1.7507	0.0261	**1.3489**	**0.0117**

Overall, our proposed method achieves better accuracy in 7 out of 10 error indexes for camera pose estimation and motion estimation of the 4 moving boxes. In particular, our method achieves 50% improvements in estimating the camera pose and motion of the swinging boxes, top-left and bottom-left. We obtain slightly higher errors when there is spinning rotational motion of the object observed, in particular the top-right swinging and rotating box, and the bottom-right rotating box. Interestingly, our proposed algorithm performs worse than Judd et al. (2018) in these cases as the algorithm is not designed for this type of motion and scenarios, but rather for the motion of relatively large objects outdoors, e.g. urban driving. We believe that this is due to using an optical flow algorithm that is not well optimised for rotating objects. The consequence of this is poor estimation of point motion and consequent degradation of the overall object tracking performance. Even with the associated performance loss for rotating objects, the benefits of dense optical flow motion estimation is clear in the other metrics.
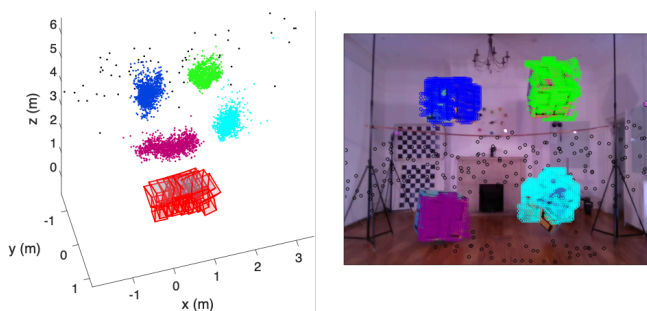

Figure 5. Qualitative results of our method on Oxford Multimotion Dataset. (Left) The 3D map including camera trajectory, static structure and tracks of dynamic points. (Right) Detected points on static background and object body. Black color corresponds to static points and features on each object are shown in a different color.

An illustrative result of the map output of our algorithm on Oxford Multimotion Dataset is shown in Fig. 5. Tracks of dynamic features on swinging boxes visually correspond to the actual motion of the boxes. This can be clearly seen in the swinging motion of the bottom-left box shown with purple color in Fig. 5.

5.4 KITTI Tracking Dataset

The KITTI Tracking Dataset (Geiger et al. (2013)) contains 21 sequences in total with ground truth information about camera and object poses. Among these sequences, some are not included in the evaluation of our system; as they contain no moving objects (static only scenes) or only contain pedestrians that are non-rigid objects, which is outside the scope of this work.
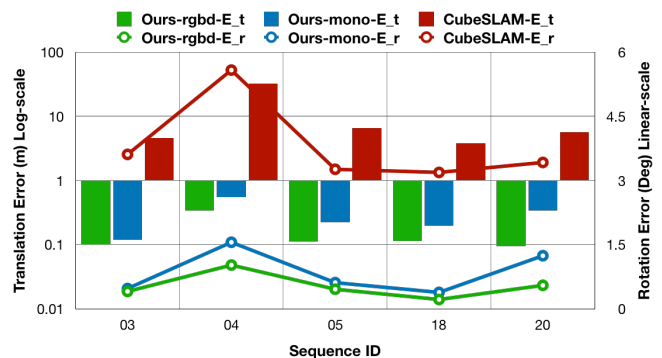

Figure 6. Accuracy of object motion estimation of our method compared to CubeSLAM (Yang and Scherer (2019)). The color bars refer to translation error that is corresponding to the left Y-axis in log-scale. The curves refer to rotation error, which corresponds to the right Y-axis in linear-scale.

5.4.1 Camera and Object Motion: Table 2 demonstrates results of both camera and object motion estimation in 9 sequences, compared to CubeSLAM (Yang and Scherer (2019)), with results of 5 sequences provided by the authors. We initially tried to evaluate CubeSLAM ourselves with the default provided parameters, however errors were much higher, and hence we only report results of the sequences provided by the authors of CubeSLAM after some correspondences. As CubeSLAM is for monocular camera, we also compute results of a learning-based monocular version of our proposed method (as mentioned in Section 4) for fair comparison.

Overall, both our proposed RGB-D and learning-based monocular methods obtain high accuracy in both camera and object motion estimation. Compared to CubeSLAM, our RGB-D version gets lower errors in camera motion, while our learning-based monocular version slightly higher. Nevertheless, both versions obtain consistently lower errors in object motion estimation. In particular, as demonstrated in Fig. 6, the translation and rotation errors in CubeSLAM are all above 3 meters and 3 degrees, with errors reaching 32 meters and 5 degrees in extreme cases respectively. However, our translation errors vary between 0.1-0.3 meters and rotation errors between 0.2-1.5 degrees in case of RGB-D, and 0.1-0.3 meters, and 0.4-3 degrees in case of learning-based monocular, which indicates that our object motion estimation achieves an order of magnitude improvements.

Table 2. Comparison versus CubeSLAM (Yang and Scherer (2019)) for camera and object motion estimation accuracy on nine sequences with moving objects drawn from the KITTI dataset. Bold numbers indicate the better result.

| | Proposed RGB-D | | | | Proposed Monocular | | | | CubeSLAM | | | |
| | Camera | | Object | | Camera | | Object | | Camera | | Object | |
seq	E_r(deg)	E_t(m)	E_r(deg)	E_t(m)	E_r(deg)	E_t(m)	E_r(deg)	E_t(m)	E_r(deg)	E_t(m)	E_r(deg)	E_t(m)
00	**0.0512**	**0.0541**	**1.0429**	**0.1061**	0.0708	0.0591	1.7242	0.3081				
01	**0.0368**	**0.1201**	**1.2504**	**0.2129**	0.1133	0.3484	1.9514	0.5389				
02	**0.0152**	**0.0427**	**1.4311**	**0.2358**	0.0285	0.0584	2.3104	0.5729				
03	**0.0377**	**0.0866**	**0.3782**	**0.1055**	0.1010	0.1628	0.4722	0.1201	0.0498	0.0929	3.6085	4.5947
04	**0.0476**	**0.1127**	**1.0211**	**0.3376**	0.1539	0.5096	1.5543	0.5503	0.0708	0.1159	5.5803	32.5379
05	**0.0246**	0.0982	**0.4482**	**0.1098**	0.0576	0.1263	0.6103	0.2245	0.0342	**0.0696**	3.2610	6.4851
06	**0.0477**	**0.0187**	**1.2156**	**0.1304**	0.0511	0.0296	3.4771	0.3187				
18	**0.0192**	0.0715	**0.1470**	**0.1131**	0.0861	0.2179	0.3818	0.1953	0.0433	**0.0510**	3.1876	3.7948
20	**0.0293**	**0.1643**	**0.5401**	**0.0944**	0.0872	0.4092	1.2395	0.3421	0.1348	0.1888	3.4206	5.6986

5.4.2 Object Tracking and Velocity:

We also demonstrate the performance of tracking dynamic objects, and show results of object speed estimation, which is an important information for autonomous driving applications. Fig. 7 illustrates results of object tracking length and object speed for some selected objects (tracked for over 20 frames) in all the tested sequences. Our system is able to track most objects for more than 80% of their occurrence in the sequence. Moreover, our estimated objects speed is always consistently close to the ground truth.
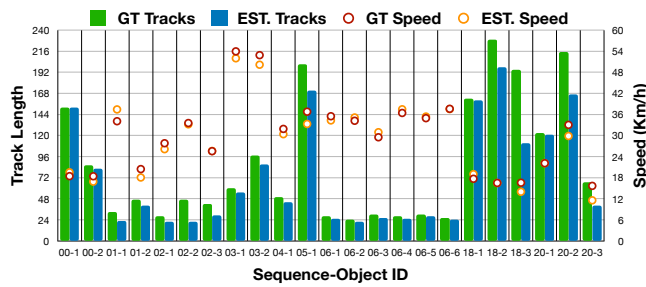

Figure 7. Tracking performance and speed estimation. Results of object tracking length and object speed for some selected objects (tracked for over 20 frames), due to limited space. The color bars represent the length of object tracks, which is corresponding to the left Y-axis. The circles represent object speeds, which is corresponding to the right Y-axis. GT refers to ground truth, and EST. refers to estimated values.

5.4.3 Qualitative Results:

Fig. 8 illustrates the output of our system for three of the KITTI sequences. The proposed system is able to output the camera poses, along with the static structure and dynamic tracks of every detected moving object in the scene in a spatio-temporal map representation.

5.5 Discussion

Apart from the extensive evaluation in Section 5.4 and 5.3, we also provide detailed experimental results to prove the effectiveness of key modules in our proposed system. Finally, the computational cost of the proposed system is discussed.

5.5.1 Robust Tracking of Points:

The graph optimisation explores the spacial and temporal information to refine the camera poses and the object motions, as well as the static and dynamic structure. This process requires robust tracking of good points in terms of both quantity and quality. This was achieved by refining the estimated optical flow jointly with the motion estimation, as discussed in Section 3.2.3. The effectiveness of joint optimisation is shown by comparing a baseline method that only optimises for the motion (Motion Only) using (9) for camera motion or (11) for object motion, and the improved method that optimises for both the motion and the optical flow (Joint) using (13) or (15). Table 3 demonstrates that the joint method obtains considerably more points that are tracked for long periods.

Table 3. The number of points tracked for more than five frames on the nine sequences of the KITTI dataset. Bold numbers indicate the better results. Underlined bold numbers indicate an order of magnitude increase in number.

| | Background | | Object | |
seq	Motion Only	Joint	Motion Only	Joint
00	1798	**_12812_**	1704	**7162**
01	237	**_5075_**	907	**4583**
02	7642	**10683**	52	**_1442_**
03	778	**_12317_**	343	**_3354_**
04	9913	**25861**	339	**_2802_**
05	713	**_11627_**	2363	**2977**
06	7898	**11048**	482	**_5934_**
18	4271	**22503**	5614	**14989**
20	9838	**49261**	9282	**13434**

Using the tracked points given by the joint estimation process leads to better estimation of both camera pose and object motion. An improvement of about 15% to 20% in translation and rotation errors was observed over the nine sequences of the KITTI dataset shown above. See Table 4.

Table 4. Average camera pose and object motion errors over the nine sequences of the KITTI dataset. Bold numbers indicate the better results.

| | Motion Only | | Joint | |
	E_r(deg)	E_t(m)	E_r(deg)	E_t(m)
Camera	0.0412	0.0987	**0.0344**	**0.0854**
Object	1.0179	0.1853	**0.8305**	**0.1606**

5.5.2 Robustness against Non-direct Occlusion:

The mask segmentation may fail in some cases, due to direct or indirect occlusions (illumination change, etc.). Thanks to the mask propagating method described in Section 4.3.3, our proposed system is able to handle mask failure cases caused by indirect occlusions. Fig. 9 demonstrates an example of tracking a white van for 80 frames, where the mask segmentation fails in 33 frames. Despite the object segmentation failure, our system is still continuously able to track the van, and estimate its speed with an average error of 2.64 km/h across the whole sequence. Speed errors in the second half of the sequence are higher due to partial direct occlusions, and increased distance to the object getting farther away from the camera.
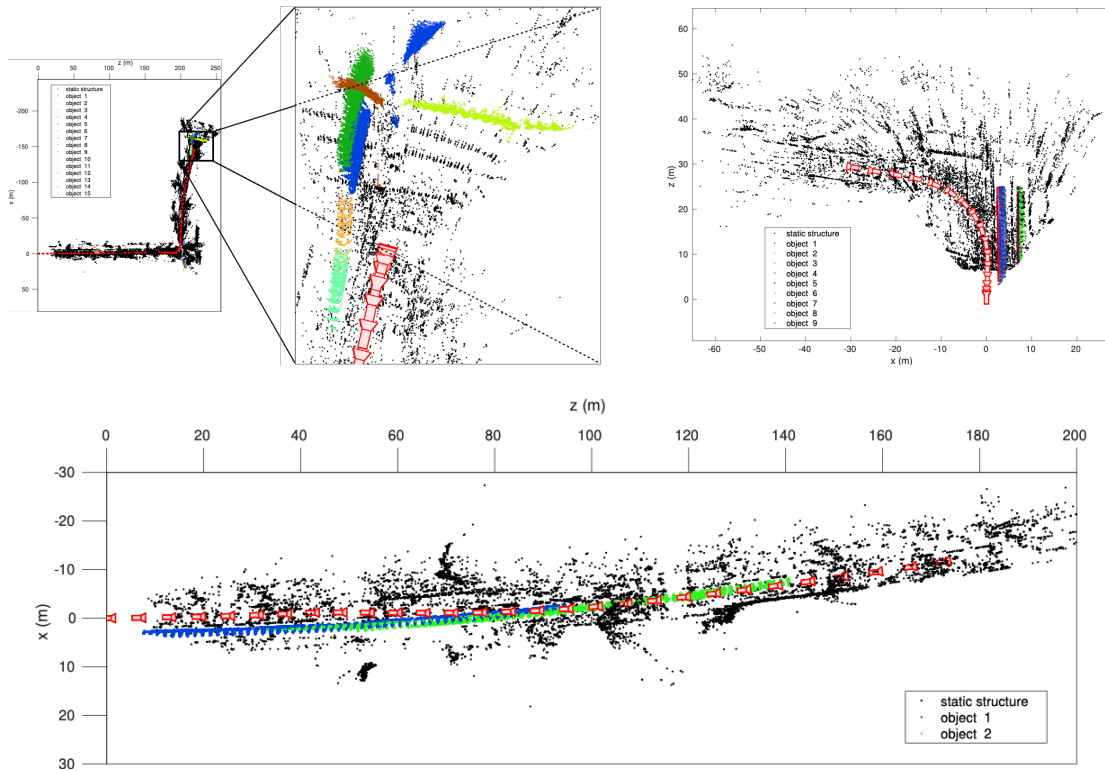
Figure 8. Illustration of system output; a dynamic map with camera poses, static background structure, and tracks of dynamic objects. Sample results of VDO-SLAM on KITTI sequences. Black represents static background, and each detected object is shown in a different colour. Top left figure represents seq.01 and a zoom-in on the intersection at the end of the sequence, top right figure represents seq.06 and bottom figure represents seq.03.
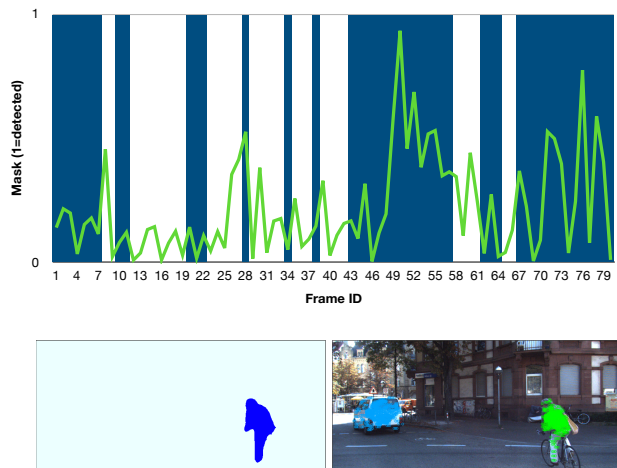

Figure 9. Robustness in tracking performance and speed estimation in case of semantic segmentation failure. An example of tracking performance and speed estimation for a white van (ground-truth average speed 20km/h) in seq.00. (Top) Blue bars represent a successful object segmentation, and green curves refer to the object speed error. (Bottom-left) An illustration of semantic segmentation failure on the van. (Bottom-right) Result of propagating the previously tracked features on the van by our system.

5.5.3 Global Refinement on Object Motion: Initial object motion estimation (in the tracking component of the system) is independent between frames, since it is purely related to the sensor measurements. As illustrated in Fig. 10, the

blue curve describes an initial object speed estimate of a wagon observed for 55 frames in sequence 03 of the KITTI tracking dataset. As seen in the figure, the speed estimation is not smooth and large errors occur towards the second half of the sequence. This is mainly caused by the increased distance to the object getting farther away from the camera, and its structure only occupying a small portion of the scene. In this case, the object motion estimation from sensor measurements solely becomes challenging and error-prone. Therefore, we formulate a factor graph and refine the motions together with the static and dynamic structure as discussed in Section 3.3. The green curve in Fig. 10 shows the object speed results after the global refinement, which becomes smoother in the first half of the sequence and is significantly improved in the second half.

5.5.4 Computational Analysis: Finally, we provide the computational analysis of our system. The experiments are carried out on an Intel Core i7 2.6 GHz laptop computer with 16 GB RAM. The object semantic segmentation and dense optical flow computation times depend on the GPU power and the CNN model complexity. Many current state-of-the-art algorithms can run in real time (Bolya et al. (2019); Hui et al. (2020)). In this paper, the semantic segmentation and optical flow results are produced off-line as input to the system. The SLAM system is implemented in C++ on CPU using a modified version of g2o as a back-end (Kümmerle et al. (2011)). We show the computational time in Table 5 for both datasets. In the local batch optimisation, the window size is set to 20 frames with an overlap of 4 frames. The time cost of every system component is averaged over all
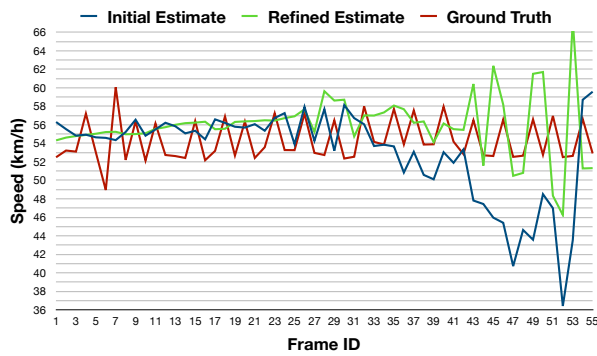
Figure 10. Global refinement effect on object motion estimation. The initial (blue) and refined (green) estimated speeds of a wagon in seq.03, travelling along a straight road, compared to the ground truth speed (red).

frames, and sequences. Overall, the tracking part of our proposed system is able to run at the frame rate of 5-8 fps depending on the number of detected moving objects, which can be improved by employing parallel implementation. The runtime of the global batch optimisation strongly depends on the amount of camera poses (number of frames), and objects (density in terms of the number of dynamic objects observed per frame) present in the scene.

Table 5. Runtime of different system components for both datasets. The time cost of every component is averaged over all frames and sequences, except for the object motion estimation that is averaged over the number of objects, and local batch optimisation that is averaged over the number of frames.

Dataset	Tasks	Runtime (mSec)
KITTI	Feature Detection	16.2550
	Camera Pose Estimation	52.6542
	Dynamic Object Tracking	11.4828
	Object Motion Estimation (avg/object)	22.9081
	Map and Mask Updating	22.1830
	Local Batch Optimisation (avg /frame)	18.2828
OMD	Feature Detection	7.5220
	Camera Pose Estimation	32.0909
	Dynamic Object Tracking	28.0535
	Object Motion Estimation (avg/object)	19.5280
	Map and Mask Updating	30.3153
	Local Batch Optimisation (avg /frame)	15.3414

6   Conclusion

In this paper, we have presented VDO-SLAM, a novel dynamic feature-based SLAM system that exploits image-based semantic information in the scene with no additional knowledge of the object pose or geometry, to achieve simultaneous localisation, mapping and tracking of dynamic objects. The system consistently shows robust and accurate results on indoor and challenging outdoor datasets, and achieves state-of-the-art performance in object motion estimation. We believe the high performance accuracy achieved in object motion estimation is due to the fact that our system is a feature-based system. Feature points remain to be the easiest to detect, track and integrate within a SLAM system, and that require the front-end to have no additional knowledge about the object model, or explicitly provide any information about its pose.

An important issue to be reduced is the computational complexity of SLAM with dynamic objects. In long-term

applications, different techniques can be applied to limit the growth of the graph (Strasdat et al. (2011); Ila et al. (2010)). In fact, history summarisation/deletion of map points pertaining to dynamic objects observed far in the past seems to be a natural step towards a long-term SLAM system in highly dynamic environments.

7   Acknowledgements

References

Agarwal S, Mierle K and Others (2012) Ceres solver. http://ceres-solver.org.

Alcantarilla PF, Yebes JJ, Almazán J and Bergasa LM (2012) On combining visual SLAM and dense scene flow to increase the robustness of localization and mapping in dynamic environments. In: *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, pp. 1290–1297.

Bescos B, Fácil JM, Civera J and Neira J (2018) DynaSLAM: Tracking, mapping, and inpainting in dynamic scenes. *IEEE Robotics and Automation Letters* 3(4): 4076–4083.

Bibby C and Reid I (2007) Simultaneous Localisation and Mapping in Dynamic Environments (SLAMIDE) with reversible data association. In: *Proceedings of Robotics: Science and Systems*.

Bolya D, Zhou C, Xiao F and Lee YJ (2019) YOLACT: real-time instance segmentation. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 9157–9166.

Butler DJ, Wulff J, Stanley GB and Black MJ (2012) A naturalistic open source movie for optical flow evaluation. In: *European conference on computer vision*. Springer, pp. 611–625.

Byravan A and Fox D (2017) SE3-Nets: Learning rigid body motion using deep neural networks. In: *IEEE International Conference on Robotics and Automation (ICRA), 2017*. IEEE, pp. 173–180.

Cadena C, Carlone L, Carrillo H, Latif Y, Scaramuzza D, Neira J, Reid I and Leonard JJ (2016) Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics* 32(6): 1309–1332.

Chirikjian GS, Mahony R, Ruan S and Trumpf J (2017) Pose changes from a different point of view. In: *Proceedings of the ASME International Design Engineering Technical Conferences (IDETC) 2017*. ASME.

Chu Q, Ouyang W, Li H, Wang X, Liu B and Yu N (2017) Online multi-object tracking using cnn-based single object tracker with spatial-temporal attention mechanism. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 4836–4845.

de la Puente P and Rodríguez-Losada D (2014) Feature based graph-SLAM in structured environments. *Autonomous Robots* 37(3): 243–260.

Dellaert F and Kaess M (2006) Square Root Sam: Simultaneous localization and mapping via square root information smoothing. *The International Journal of Robotics Research* 25(12): 1181–1203.

Eigen D, Puhrsch C and Fergus R (2014) Depth map prediction from a single image using a multi-scale deep network. In: *Advances in neural information processing systems*. pp. 2366–2374.

Elhamifar E and Vidal R (2009) Sparse subspace clustering. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, pp. 2790–2797.

Gálvez-López D, Salas M, Tardós JD and Montiel J (2016) Real-time monocular object SLAM. *Robotics and Autonomous Systems* 75: 435–449.

Geiger A, Lenz P, Stiller C and Urtasun R (2013) Vision meets robotics: The KITTI dataset. *The International Journal of Robotics Research* 32(11): 1231–1237.

Geiger A, Lenz P and Urtasun R (2012) Are we ready for autonomous driving? the KITTI vision benchmark suite. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Girshick R, Radosavovic I, Gkioxari G, Dollár P and He K (2018) Detectron. https://github.com/facebookresearch/detectron.

Godard C, Mac Aodha O, Firman M and Brostow GJ (2019) Digging into self-supervised monocular depth estimation. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 3828–3838.

Hahnel D, Schulz D and Burgard W (2002) Map building with mobile robots in populated environments. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2002.*, volume 1. IEEE, pp. 496–501.

Hahnel D, Triebel R, Burgard W and Thrun S (2003) Map building with mobile robots in dynamic environments. In: *IEEE International Conference on Robotics and Automation, 2003. Proceedings. ICRA'03*, volume 2. IEEE, pp. 1557–1563.

He K, Gkioxari G, Dollár P and Girshick R (2017) Mask R-CNN. In: *IEEE International Conference on Computer Vision (ICCV), 2017*. IEEE, pp. 2980–2988.

Henein M, Abello M, Ila V, and Mahony R (2017) Exploring The Effect of Meta-Structural Information on the Global Consistency of SLAM. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems 2017*. The Australian National University.

Henein M, Zhang J, Mahony R and Ila V (2020) Dynamic SLAM: The need for speed. *2020 IEEE International Conference on Robotics and Automation (ICRA). To appear.* .

Hsiao M, Westman E, Zhang G and Kaess M (2017) Keyframe-based dense planar SLAM. In: *IEEE International Conference on Robotics and Automation (ICRA), 2017*. IEEE, pp. 5110–5117.

Huber PJ (1992) Robust estimation of a location parameter. In: *Breakthroughs in statistics*. Springer, pp. 492–518.

Hui TW, Tang X and Loy CC (2020) A Lightweight Optical Flow CNN - Revisiting Data Fidelity and Regularization. URL http://mmlab.ie.cuhk.edu.hk/projects/LiteFlowNet/.

Ila V, Polok L, Šolony M and Svoboda P (2017) SLAM++-A highly efficient and temporally scalable incremental SLAM framework. *International Journal of Robotics Research* Online First(0): 1–21. DOI:10.1177/0278364917691110.

Ila V, Porta JM and Andrade-Cetto J (2010) Information-based compact pose SLAM. *IEEE Transactions on Robotics* 26(1): 78–93.

Judd KM and Gammell JD (2019) The Oxford Multimotion Dataset: Multiple SE(3) Motions with Ground Truth. *IEEE Robotics and Automation Letters* 4(2): 800–807.

Judd KM, Gammell JD and Newman P (2018) Multimotion visual odometry (MVO): Simultaneous estimation of camera and third-party motions. In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pp. 3949–3956.

Kaess M (2015) Simultaneous Localization and Mapping with Infinite Planes. In: *IEEE International Conference on Robotics and Automation (ICRA), 2015*. IEEE, pp. 4605–4611.

Kaess M, Johannsson H, Roberts R, Ila V, Leonard JJ and Dellaert F (2011) iSAM2: Incremental smoothing and mapping using the bayes tree. *The International Journal of Robotics Research* : 0278364911430419.

Ke T and Roumeliotis SI (2017) An efficient algebraic solution to the perspective-three-point problem. In: *CVPR*.

Khan Z, Balch T and Dellaert F (2005) Mcmc-based particle filtering for tracking a variable number of interacting targets. *IEEE transactions on pattern analysis and machine intelligence* 27(11): 1805–1819.

Kim C, Li F and Rehg JM (2018) Multi-object tracking with neural gating using bilinear lstm. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 200–215.

Kümmerle R, Grisetti G, Strasdat H, Konolige K and Burgard W (2011) g2o: A general framework for graph optimization. In: *IEEE International Conference on Robotics and Automation (ICRA), 2011*. IEEE, pp. 3607–3613.

Kundu A, Krishna KM and Jawahar C (2011) Realtime multibody visual SLAM with a smoothly moving monocular camera. In: *IEEE International Conference on Computer Vision (ICCV), 2011*. IEEE, pp. 2080–2087.

Lai T, Wang H, Yan Y, Chin TJ and Zhao WL (2016) Motion segmentation via a sparsity constraint. *IEEE Transactions on Intelligent Transportation Systems* 18(4): 973–983.

Lin TY, Maire M, Belongie S, Hays J, Perona P, Ramanan D, Dollár P and Zitnick CL (2014) Microsoft COCO: Common objects in context. In: *European conference on computer vision*. Springer, pp. 740–755.

Lv Z, Kim K, Troccoli A, Rehg J and Kautz J (2018) Learning rigidity in dynamic scenes with a moving camera for 3d motion field estimation. In: *European Conference on Computer Vision*.

Mayer N, Ilg E, Hausser P, Fischer P, Cremers D, Dosovitskiy A and Brox T (2016) 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.

Milan A, Leal-Taixé L, Reid I, Roth S and Schindler K (2016) MOT16: A benchmark for multi-object tracking. *arXiv:1603.00831 [cs]* URL http://arxiv.org/abs/1603.00831. ArXiv: 1603.00831.

Milan A, Rezatofighi SH, Dick A, Reid I and Schindler K (2017) Online multi-target tracking using recurrent neural networks. In: *Thirty-First AAAI Conference on Artificial Intelligence*.

Miller I and Campbell M (2007) Rao-blackwellized particle filtering for mapping dynamic environments. In: *IEEE International Conference on Robotics and Automation, 2007*. IEEE, pp. 3862–3869.

Mu B, Liu SY, Paull L, Leonard J and How JP (2016) SLAM with objects using a nonparametric pose graph. In: *IEEE/RSJ*

International Conference on Intelligent Robots and Systems (IROS), 2016. IEEE, pp. 4602–4609.

Nicholson L, Milford M and Sünderhauf N (2018) QuadricSLAM: Dual Quadrics as SLAM Landmarks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. pp. 313–314.

Nistér D, Naroditsky O and Bergen J (2004) Visual odometry. In: *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1. IEEE, pp. I–I.

Polok L, Ila V, Solony M, Smrz P and Zemcik P (2013) Incremental Block Cholesky Factorization for Nonlinear Least Squares in Robotics. In: *Proceedings of Robotics: Science and Systems*. Berlin, Germany. DOI:10.15607/RSS.2013.IX.042.

Rao SR, Yang AY, Sastry SS and Ma Y (2010) Robust algebraic segmentation of mixed rigid-body and planar motions from two views. *International journal of computer vision* 88(3): 425–446.

Reddy ND, Singhal P, Chari V and Krishna KM (2015) Dynamic body VSLAM with semantic constraints. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015*. IEEE, pp. 1897–1904.

Rogers JG, Trevor AJ, Nieto-Granda C and Christensen HI (2010) SLAM with expectation maximization for moveable object tracking. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2010*. IEEE, pp. 2077–2082.

Rosten E and Drummond T (2006) Machine learning for high-speed corner detection. In: *European conference on computer vision*. Springer, pp. 430–443.

Rubino C, Del Bue A and Chin TJ (2018) Practical motion segmentation for urban street view scenes. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 1879–1886.

Rublee E, Rabaud V, Konolige K and Bradski G (2011) ORB: An efficient alternative to SIFT or SURF. In: *2011 International conference on computer vision*. Ieee, pp. 2564–2571.

Salas-Moreno RF, Newcombe RA, Strasdat H, Kelly PH and Davison AJ (2013) SLAM++: Simultaneous Localisation and Mapping at the Level of Objects. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2013*. IEEE, pp. 1352–1359.

Strasdat H, Davison AJ, Montiel JM and Konolige K (2011) Double window optimisation for constant time visual SLAM. In: *IEEE International Conference on Computer Vision (ICCV), 2011*. IEEE, pp. 2352–2359.

Sun D, Yang X, Liu MY and Kautz J (2018) PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume.

Tan W, Liu H, Dong Z, Zhang G and Bao H (2013) Robust monocular SLAM in dynamic environments. In: *Mixed and Augmented Reality (ISMAR), 2013 IEEE International Symposium on*. IEEE, pp. 209–218.

Vidal R and Hartley R (2004) Motion segmentation with missing data using PowerFactorization and GPCA. In: *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 2. IEEE, pp. II–II.

Wang CC, Thorpe C and Thrun S (2003) Online simultaneous localization and mapping with detection and tracking of moving objects: Theory and results from a ground vehicle in crowded urban areas. In: *IEEE International Conference on Robotics and Automation, 2003. Proceedings. ICRA'03*, volume 1. IEEE, pp. 842–849.

Wang CC, Thorpe C, Thrun S, Hebert M and Durrant-Whyte H (2007) Simultaneous localization, mapping and moving object tracking. *The International Journal of Robotics Research* 26(9): 889–916.

Wohlhart P and Lepetit V (2015) Learning descriptors for object recognition and 3d pose estimation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3109–3118.

Wolf DF and Sukhatme GS (2005) Mobile robot simultaneous localization and mapping in dynamic environments. *Autonomous Robots* 19(1): 53–65.

Xu X, Fah Cheong L and Li Z (2018) Motion segmentation by exploiting complementary geometric models. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2859–2867.

Yamaguchi K, McAllester D and Urtasun R (2014) Efficient joint segmentation, occlusion labeling, stereo and flow estimation. In: *European Conference on Computer Vision*. Springer, pp. 756–771.

Yang S and Scherer S (2019) CubeSLAM: Monocular 3-D Object SLAM. *IEEE Transactions on Robotics* 35(4): 925–938.

Zhao H, Chiba M, Shibasaki R, Shao X, Cui J and Zha H (2008) SLAM in a dynamic large outdoor environment using a laser scanner. In: *IEEE International Conference on Robotics and Automation, 2008. ICRA 2008*. IEEE, pp. 1455–1462.

Zhou X, Yu H, Liu H and Li Y (2015) Tracking multiple video targets with an improved gm-phd tracker. *Sensors* 15(12): 30240–30260.

Zou D and Tan P (2013) CoSLAM: Collaborative visual SLAM in dynamic environments. *IEEE transactions on pattern analysis and machine intelligence* 35(2): 354–366.