



HAL
open science

TwistSLAM: Constrained SLAM in Dynamic Environment

Mathieu Gonzalez, Eric Marchand, Amine Kacete, Jerome Royan

► **To cite this version:**

Mathieu Gonzalez, Eric Marchand, Amine Kacete, Jerome Royan. TwistSLAM: Constrained SLAM in Dynamic Environment. IEEE Robotics and Automation Letters, 2022, 7 (3), pp.6846-6853. hal-03676527

HAL Id: hal-03676527

<https://inria.hal.science/hal-03676527v1>

Submitted on 24 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

TwistSLAM: Constrained SLAM in Dynamic Environment

Mathieu Gonzalez¹, Eric Marchand², Amine Kacete¹ and Jerome Royan¹

Abstract— Classical visual simultaneous localization and mapping (SLAM) algorithms usually assume the environment to be rigid. This assumption limits the applicability of those algorithms as they are unable to accurately estimate the camera poses and world structure in real life scenes containing moving objects (e.g. cars, bikes, pedestrians, etc.). To tackle this issue, we propose TwistSLAM: a semantic, dynamic and stereo SLAM system that can track dynamic objects in the environment. Our algorithm creates clusters of points according to their semantic class. Thanks to the definition of inter-cluster constraints modeled by mechanical joints (function of the semantic class), a novel constrained bundle adjustment is then able to jointly estimate both poses and velocities of moving objects along with the classical world structure and camera trajectory. We evaluate our approach on several sequences from the public KITTI dataset and demonstrate quantitatively that it improves camera and object tracking compared to state-of-the-art approaches.

Index Terms—SLAM, Localization, Mapping

I. INTRODUCTION

VISUAL Simultaneous Localization And Mapping (SLAM) is an important problem for robotics that has been heavily studied in the past decade. Its goal is to estimate the pose of a camera moving in a scene while building a map of the environment. Some algorithms such as [1], [2] can efficiently solve this problem, however they rely on the static scene assumption. This hypothesis assumes that the world is a single rigid body and thus that no object can move within it. This assumption, which is rarely met in most real world scenes, as they can contain moving objects (e.g. cars, bikes, pedestrians for urban scenes of the KITTI dataset), limits the scenarios in which a SLAM algorithm can be used. Classical SLAM systems such as [1] try to alleviate this assumption using robust estimators, allowing them to flag moving parts as outliers. However as soon as the number of moving points is too important, the estimated camera pose accuracy decreases. This makes this approach unsuitable for some scenes (e.g. crowded or urban scenes). Some systems [3], [4] have been proposed to detect and mask out dynamic objects in images, thus making the static scene assumption valid. However some recent approaches

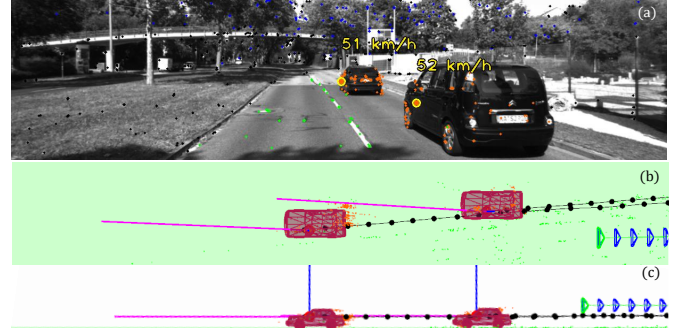


Fig. 1. Our approach allows us to track objects in the scene such as cars. Here we can see: (a) the frame with semantic points and tracked clusters (orange cars) with their estimated speed. (b) a map top view with tracked clusters (orange cars), clusters trajectories (black spheres and lines), clusters twists (blue and purple lines), road points and plane (in green) and camera trajectory (green and blue frustums). (c) Map side view. The rotation part of the twists (blue lines) is perpendicular to the road plane, the translation part (purple lines) is parallel to the plane.

[5], [6], [7], [8] argue that moving objects represent valuable information that can be necessary for some applications. Most recent approaches trying to solve both SLAM and object tracking have used semantics as an additional source of information. Semantic knowledge can indeed be beneficial to SLAM [9], [10] as it contains information about the class dynamicity [3] which is higher level information than simple 3D points.

In this paper we present a stereo SLAM system called TwistSLAM, as we estimate objects' twists to track them and consider that objects are linked to each other through mechanical joints, similarly to joints linking different parts of a robot. An illustration of our algorithm is visible in figure 1: the camera pose is estimated simultaneously with all moving objects in the scene and the map structure (here the plane of the road) constrains the movement of objects. Our approach is based on ORB-SLAM2 [1] and S³LAM [11]. In our work we use semantic information to build a map of clusters corresponding to objects in the scene. The clustering of the scene allows us to estimate the pose of the camera using static clusters only such as *road* or *house*. The other clusters that can be dynamic are tracked and their pose is updated in the map through the estimation of twists that represent their velocity. Most SLAM systems that can track dynamic objects directly estimate their pose through the minimization of a reprojection error function [5], [8] or with 3D points registration [7]. Doing so, the estimated pose of an object has 6 degrees of freedom. This does not correspond to reality. For example a car has only 3 degrees of freedom, 2 translations in the road plane and 1 rotation around its normal, hence its pose should be constrained. Our goal is thus to remove degrees

Manuscript received: February, 24, 2022; Revised April, 28, 2022; Accepted May, 19, 2022.

This paper was recommended for publication by Editor Javier Civera upon evaluation of the Associate Editor and Reviewers' comments.

¹ Mathieu Gonzalez, Amine Kacete and Jerome Royan are with the Institute of Research and Technology b<>com, Rennes, France, {mathieu.gonzalez, amine.kacete, jerome.royan}@b-com.com

² Eric Marchand is with Univ Rennes, Inria, IRISA, CNRS, Rennes, France, Eric.Marchand@irisa.fr

Digital Object Identifier (DOI): see top of this page.

of freedom corresponding to physically unfeasible movements. To do so we choose to represent those constraints as mechanical joints which makes our approach highly generic. A mechanical joint between clusters constrains the estimated twist of a dynamic cluster by blocking some of its degrees of freedom thus reducing the effect of noise on the estimation. Once an object twist has been estimated it can be used to update the object pose which enables object tracking. The object poses can then be tightly refined with camera poses and 3D points within a bundle adjustment that also applies mechanical joints constraints.

Our contributions presented in this paper are:

- A semantic SLAM system that can robustly estimate the pose of a camera in static as well as dynamic scenes.
- A stereo SLAM framework that can track multiple moving objects in the scene.
- A new formulation for both the tracking and bundle adjustment that takes into account the characteristics of mechanical joints between objects in the scene to constrain their movements.

We also evaluate our approach on several sequences from the public KITTI dataset and compare our approach quantitatively with respect to ORB-SLAM2 [1], DynaSLAM [3], VDO-SLAM [6] and DynaSLAM II [5], which demonstrates the benefits of our method in terms of object poses and velocities and camera pose estimation accuracy. The rest of the paper is described as follows. First we describe related work on dynamic classical and semantic SLAM. Then we give an overview of mathematical concepts that will be used in this paper. Following, we describe our approach to build a semantic map of clusters, estimate the pose of a camera in a dynamic scene, track moving clusters within the scene while using joint constraints to improve object tracking and refine all estimations with a bundle adjustment. Finally we demonstrate the benefits of our approach on multiple sequences from a public dataset.

II. RELATED WORK: DYNAMIC SLAM

In this section we first present some classical [1] and dynamic SLAM systems [12] that remove dynamic outliers. Then we focus on semantic SLAM systems that tackle the problem of dynamic objects by masking out [3], [4] or tracking [7], [5], [8], [6], [13] objects in the scene. Our approach belongs to this last class of systems. For a more in-depth survey of SLAM we refer the reader to [9], [14].

ORB-SLAM2 [1], which follows the work of [15], [16] uses three parallel threads to estimate the pose of a camera through relocalization or tracking, build a map, refine both using a bundle adjustment [17] and close looping trajectories. ORB-SLAM2 manages to mitigate the influence of dynamic objects through the use of a RANSAC scheme and a robust cost function based estimation. To tackle this problem some methods propose to roughly estimate the pose of the camera and robustly find outliers in the scene or in the image. Outliers are then removed or downweighted and the camera pose is refined. For example [12] proposes a direct approach based on the alignment of depth edges using an ICP scheme. For each point they robustly estimate a staticity confidence score which downweights dynamic objects and an intensity assisted ICP robustly refines the pose using those

weights. However those approaches fail when the amount of dynamic parts in the image is too high and are not able to track dynamic objects. DS-SLAM [4] applies a geometrical moving consistency check on segmented areas. This score allows them to know which areas correspond to moving objects, which are then discarded for robust camera pose estimation and mapping. DynaSLAM [3] uses semantic information to segment a priori moving objects, which are not used for tracking and mapping. The segmentation is refined using depth information. This approach improves camera localization in dynamic scenes but deteriorates it when a priori moving objects are in reality static such as parked cars. Those approaches, similarly to the previous ones, are not able to track dynamic objects. MaskFusion [7] is one of the first semantic dynamic SLAM that can track objects. Inspired from [18] it makes use of 2D instance segmentation to detect objects in the scene and tracks them using both photometric and geometric information from an RGB-D camera. DynaSLAM II [5] uses semantic information to detect objects. Object 3D points are represented in the object reference frame and used to estimate the object pose at all time by minimizing their reprojection error. ClusterVO [8] is similar to [5], but they consider object detection (i.e. 2D bounding boxes) as input which is much faster to infer than dense masks. They also apply a cleaning procedure to improve dynamic keypoints matching and make sure that 2D points do not come from the background of the bounding box. VDO-SLAM [6] proposes to use optical flow to track features extracted more densely than other systems, which allows them to obtain a more precise object pose estimation. Furthermore the optical flow and the object and camera motions are tightly refined. CubeSLAM [13] is different from the previously cited papers as it is an object based SLAM. A 3D bounding box is fitted for each object detected by a CNN (such as [19]), which allows them to know the object 6 DoF pose and dimensions for each image. The bounding boxes poses are then optimized along with the camera poses in a single BA, similarly to [5].

III. BACKGROUND ON TWISTS AND HOMOGENEOUS TRANSFORMATIONS

In this section we present the mathematical concepts and notations that we use in this paper. A rigid object o , which can move in the 3D scene, can be associated with its coordinate frame \mathcal{F}_o . Its pose relative to the world coordinate frame \mathcal{F}_w can be represented by the homogeneous matrix in the i^{th} frame

$${}^w\mathbf{T}_{o_i} = \begin{pmatrix} {}^w\mathbf{R}_{o_i} & {}^w\mathbf{t}_{o_i} \\ 0 & 1 \end{pmatrix} \in \text{SE}(3) \quad (1)$$

where $\text{SE}(3)$ denotes the special euclidean group. This matrix maps points expressed in the object coordinate frame and denoted ${}^o\mathbf{X}$ to points expressed in the world coordinate frame and denoted ${}^w\mathbf{X}$, according to the following equation: ${}^w\mathbf{X} = {}^w\mathbf{T}_o {}^o\mathbf{X}$. The velocity of a moving object can be represented using a twist ξ defined as

$$\xi = (v_x \ v_y \ v_z \ \omega_x \ \omega_y \ \omega_z)^\top = (\mathbf{v} \ \boldsymbol{\omega})^\top \in \mathbb{R}^6 \quad (2)$$

where the first 3 components $\mathbf{v} = (v_x, v_y, v_z)^\top \in \mathbb{R}^3$ denote the translational velocity and the other components $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)^\top \in \mathbb{R}^3$ represent the rotational velocity. There

exists a matrix representation for twists that can be obtained using the operator $[\cdot]_{\wedge}$ defined by:

$$[\boldsymbol{\xi}]_{\wedge} = \begin{pmatrix} [\boldsymbol{\omega}]_{\times} & \mathbf{v} \\ 0 & 0 \end{pmatrix} \in \mathfrak{se}(3) \quad (3)$$

where $[\cdot]_{\times}$ is the skew-symmetric operator defined such as for $\mathbf{a} = (a_x, a_y, a_z)^{\top} \in \mathbb{R}^3$:

$$[\mathbf{a}]_{\times} = \begin{pmatrix} 0 & -a_z & a_y \\ a_z & 0 & -a_x \\ -a_y & a_x & 0 \end{pmatrix} \quad (4)$$

and $\mathfrak{se}(3)$ is the Lie algebra associated to $\text{SE}(3)$. We denote ${}^w \boldsymbol{\xi}_{o_i}$ the twist corresponding to the velocity of the object o at frame i expressed in the world coordinate frame. Similarly, ${}^{o_i} \boldsymbol{\xi}_{o_i}$ is the velocity of the object o at frame i expressed in its own coordinate frame. As velocities can be integrated over time to obtain new positions, there exists a mapping from $\mathfrak{se}(3)$ to $\text{SE}(3)$ called the exponential map and denoted:

$$\exp : [\boldsymbol{\xi} \delta t]_{\wedge} \in \mathfrak{se}(3) \longrightarrow \mathbf{T} \in \text{SE}(3) \quad (5)$$

where δt is the time interval duration [20].

Using the exponential map we can recover the pose of the object o moving according to the twist ${}^w \boldsymbol{\xi}_{o_i}$ from its initial pose at frame i , ${}^w \mathbf{T}_{o_i}$ to its next pose at frame $i+1$, ${}^w \mathbf{T}_{o_{i+1}}$ using the following formula:

$${}^w \mathbf{T}_{o_{i+1}} = \exp({}^w \boldsymbol{\xi}_{o_i} \delta t_i) {}^w \mathbf{T}_{o_i} = {}^w \mathbf{T}_{o_i} \exp({}^{o_i} \boldsymbol{\xi}_{o_i} \delta t_i) \quad (6)$$

Note that the choice of coordinate frame matters, it can be useful to define an operator to change the coordinate frame of a twist. Such operator is called the adjoint map ${}^w \mathbf{V}_{o_i} \in \mathbb{R}^{6 \times 6}$ and links twists in different coordinate frames according to:

$${}^w \boldsymbol{\xi}_{o_i} = {}^w \mathbf{V}_{o_i} {}^{o_i} \boldsymbol{\xi}_{o_i} \quad (7)$$

The adjoint map can be computed using the relative pose ${}^w \mathbf{T}_{o_i}$ between \mathcal{F}_{o_i} and \mathcal{F}_w :

$${}^w \mathbf{V}_{o_i} = \begin{pmatrix} {}^w \mathbf{R}_{o_i} & [{}^w \mathbf{t}_{o_i}]_{\times} {}^w \mathbf{R}_{o_i} \\ 0 & {}^w \mathbf{R}_{o_i} \end{pmatrix} \quad (8)$$

For simplicity we will consider in the remainder of this paper that $\delta t = 1$ without loss of generality.

IV. TWISTSLAM: CONSTRAINED SLAM IN DYNAMIC ENVIRONMENT

In this section we present our approach. Our general idea is to represent the world as a graph of semantic clusters, which is similar to a scene graph and can be seen in figure 2. The vertices of the graph correspond to objects in the scene and the edges to physical links that exist between objects. Our goal is to estimate the pose of the camera and the pose of every moving object while using mechanical joints between objects to improve those estimations. For example both clusters *car* in our graph are linked to the road with a planar constraint that allows only 3 degrees of freedom: a rotation around the normal of the plane and 2 translations within the plane. Such simple representations allow us to be highly generic as, for a given semantic class, we only need to define its static parent and the type of mechanical joint.




Fig. 2. Example of semantic graph: dynamic clusters are linked to static parent clusters with mechanical joints such as *planar* or *revolute*.

The pipeline of our approach is presented in figure 3. Using semantic information we create clusters of points corresponding to objects in the scene. Then, we use static semantic clusters (e.g. road, floor, house) to robustly track the camera, even in dynamic scenes. Next, we match keypoints corresponding to dynamic objects (e.g., car, bike,...) to either track them or triangulate new 3D points using stereo information. All poses estimations from the camera and the objects are then refined with static and dynamic 3D points in a bundle adjustment process.

The main novelty of our approach comes from the fact that we optimize the velocities of dynamic objects rather than their pose and constrain the velocities according to mechanical joints between objects. This approach is highly generic as we only need to define a handful of joints (that correspond to normalized joints in mechanics) and a list of semantic classes pairs for each joint (e.g. the wall-door joint corresponds to a revolute joint, the car-road joint corresponds to a planar joint). As we will latter show, it allows us to remove displacements along directions that are not physically possible (e.g. a car translating vertically). This allows us to obtain a more precise estimation of the dynamic object poses.

A. Creating Clusters from panoptic segmentation

Most recent semantic dynamic SLAM systems use either an object detection or an instance segmentation algorithm. Working in the continuity of S^3SLAM [11] we chose to estimate the panoptic segmentation of images, obtained using [21]. This allows us to know the semantic class of each pixel in the image and to give a unique id to each object. Similarly to [11] we fuse multiple 2D observations of a single 3D point to obtain its class and id. Doing so we obtain a semantic map, which allows us to create a set of K clusters $\mathcal{O} = \{O_k, k \in [1, K]\}$. A cluster is a set of 3D points corresponding to a single object in the scene. Points are grouped according to their class and instance id. The set of clusters can be expressed as the set of a priori static clusters \mathcal{S} (such as *road, building, ...*) and the set of a priori dynamic clusters \mathcal{D} (such as *car, bike, human, bus, ...*). As static clusters are fixed, we represent their 3D points $\{{}^w \mathbf{X}\}$ in the world frame. In contrast, each dynamic cluster contains a set of 3D points $\{{}^o \mathbf{X}\}$ expressed in the object coordinate frame, a set of poses $\{{}^o \mathbf{T}_w\}$ and a set of twists $\{{}^w \boldsymbol{\xi}_o\}$ representing the cluster trajectory and velocity through time. For simplicity in the remainder of this paper we will omit the object index k as its use is straightforward.




Fig. 3. The pipeline of our approach: static keypoints are extracted from a stereo image (a) and used for camera tracking (d), dynamic keypoints are extracted from bounding boxes within the stereo images (b) and matched using optical flow (c) with the previous frame to track dynamic objects (e). The keyframe is then segmented (f) to create new semantic map points and clusters. Finally the object and camera poses are jointly refined with the dynamic and static map points in a BA (g).

B. Clusters geometry

Our goal is to constrain the velocity of moving clusters according to mechanical joints. To do so we need to estimate the pose of those joints. We propose to do this using the estimated geometry of some clusters. We chose to consider only planar clusters, which allows our approach to be highly generic as planes are common in man-made environment. For clusters corresponding to a priori chosen classes (such as the road or the facade of a building) we estimate a 3D plane, represented by $\pi = (a, b, c, d)^T$ with $\|\pi\|^2 = 1$, using only its 3D points $\{^w\mathbf{X}\}$. The plane follows the following equation: $\pi^T {}^w\mathbf{X} = 0$ and can be estimated using an SVD. To make it robust to outliers we use a RANSAC scheme.

C. Dynamic SLAM

As we do not know which dynamic objects in the scene are really moving we chose to estimate the camera pose using only static objects. Using points from static clusters we minimize the following cost function

$$E({}^{c_i}\mathbf{T}_w) = \sum_{j \in \mathcal{S}} \rho(\|{}^i\mathbf{x}_j - p({}^{c_i}\mathbf{T}_w, {}^w\mathbf{X}_j)\|_{\Sigma_{i,j}^{-1}}) \quad (9)$$

where ${}^i\mathbf{x}_j$ is the 2D keypoint corresponding to the observation of ${}^w\mathbf{X}_j$ in the i^{th} frame, p is the pinhole camera projection function, ρ is a robust cost function (in our case Huber) [22] and $\Sigma_{i,j}$ is the covariance matrix of the reprojection error. Doing so the estimated camera pose does not take into account potentially moving objects, hence it is robust in dynamic scenes. However the estimation can be deteriorated in scenes that contain many potentially moving objects that are in reality static, like for example parked cars. To solve this problem, we chose to estimate the pose of all moving objects and integrate them in the bundle adjustment, so that the velocity of static objects is close to 0 and their points act as static points.

D. Dynamic data association and keypoints

Dynamic data association is a challenging problem for two reasons: first the combination of the camera and the object

movements can produce large displacements in the image space thus needing a large radius search for keypoints matching. Second a large movement can cause an important visual variation of the object in the image (e.g. due to luminosity changes on the object or to viewpoint changes) which makes the matching process more difficult. To overcome those challenges we propose to use the optical flow estimation produced by a CNN [23] to have a good estimate of the keypoints location and reduce the search radius, thus reducing both search time and the probability of false matches.

One problem of object tracking compared to classical SLAM is that dynamic objects usually occupy a small portion of the image, leading to too few object points to obtain a precise estimation [5], [6]. To solve this problem we force the keypoint extraction process to keep more keypoints from areas defined by objects bounding boxes. The keypoints are then used either to create new 3D points with stereo triangulation, which are added to existing clusters or used to create new clusters, or used to track the existing cluster.

E. Mechanical joints as inter-cluster constraints

Using matches found by the data association process we seek to estimate the pose of dynamic objects in the scene. Our assumption in this work is that many moving clusters can be represented as being linked to a static parent cluster with a specific mechanical joint. There exist 12 normalized joints (ISO 3952) that can be associated with the degrees of freedom they have. For example the planar joint has 3 degrees of freedom: 2 translations in the plane and 1 rotation around its normal, this joint can represent the displacement of a car relative to its static parent, the road. Another example is the revolute joint which has a single degree of freedom corresponding to the rotation around a single axis. In this case the static parent cluster is the wall, the moving cluster is the door and its only possible movements are rotations around the axis of the joint (corresponding to the hinge).

To easily model all types of joints, similarly to [24], we propose to decompose the space of twist as the sum of two orthogonal spaces:

$$\mathbb{R}^6 = \mathbb{F}_l + \mathbb{F}_l^\top \quad (10)$$

where \mathbb{F}_l (which stands for freedom) is the space of twists allowed by the mechanical joint l with coordinate frame \mathcal{F}_l . In the case of a planar joint with axis z , \mathbb{F}_l is defined as:

$$\mathbb{F}_l = \text{Span}((1 \ 1 \ 0 \ 0 \ 0 \ 1)^\top) \quad (11)$$

where Span is the linear span [25]. In general we note:

$$\mathbb{F}_l = \text{Span}(\mathbf{A}_l) \quad (12)$$

where \mathbf{A}_l is a basis of \mathbb{F}_l . To make the displacement of an object physically accurate, its twists have to lie within the \mathbb{F}_l space. To do so we project the twist from its original space to \mathbb{F}_l . This is straightforward as \mathbb{R}^6 is Euclidean, the operation projector is a 6×6 matrix defined as:

$$\mathbf{\Pi}_l = \mathbf{A}_l(\mathbf{A}_l^\top \mathbf{A}_l)^{-1} \mathbf{A}_l^\top \quad (13)$$

In the example of a planar joint, it is easy to compute that:

$$\mathbf{\Pi}_l = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \quad (14)$$

in that case, a general twist can be projected such that:

$$\mathbf{\Pi}_l \boldsymbol{\xi} = (v_x \ v_y \ 0 \ 0 \ 0 \ \omega_z)^\top \quad (15)$$

As we can see the only remaining degrees of freedom of the twist are coherent with the joint. Using this new constraint, we can modify the reprojection equation:

$${}^i \mathbf{x}_j = p({}^{c_i} \mathbf{T}_w \exp(\mathbf{\Pi}_l^w \boldsymbol{\xi}_{o_i}) {}^w \mathbf{T}_{o_{i-1}}, {}^o \mathbf{X}_j) \quad (16)$$

However this equation is only true if the twist is expressed in the joint coordinate frame, yet according to (6), it is naturally expressed either in the world or in the object coordinate frame. To change the coordinate frame of a twist we can use the adjoint map defined in (8). Hence the reprojection equation of the j^{th} point in frame i becomes:

$${}^i \mathbf{x}_j = p({}^{c_i} \mathbf{T}_w \exp({}^w \mathbf{V}_l \mathbf{\Pi}_l {}^l \mathbf{V}_w {}^w \boldsymbol{\xi}_{o_i}) {}^w \mathbf{T}_{o_{i-1}}, {}^o \mathbf{X}_j) \quad (17)$$

In the remainder of this paper we will note $\mathbf{\Pi} = {}^w \mathbf{V}_l \mathbf{\Pi}_l {}^l \mathbf{V}_w$ for simplicity. This equation takes a 3D point in the object frame, transforms it in the world frame using the previous object pose and multiplies it by the exponential of the current twist to get its current position. The twist is expressed in the joint coordinate frame with the adjoint map, projected using $\mathbf{\Pi}_l$ to keep only the relevant components and expressed again in the world frame with the inverse adjoint map. Doing this we obtain a 3D point in the world frame for frame i with a transformation that perfectly respects the mechanical joint. We then apply the camera pose to obtain the point in the camera coordinate frame, which allows us to project it in the image.

Using the reprojection function we can estimate the twist

corresponding to the transformation of a set of object points between frame $i - 1$ and i by minimizing the following error:

$$E({}^{o_i} \boldsymbol{\xi}_w) = \sum_j \rho(\|{}^i \mathbf{x}_j - p({}^{c_i} \mathbf{T}_w \exp(\mathbf{\Pi}^w \boldsymbol{\xi}_{o_i}) {}^w \mathbf{T}_{o_{i-1}}, {}^o \mathbf{X}_j)\|_{\Sigma_{i,j}^{-1}}) \quad (18)$$

where ρ is the Huber robust estimator [22] and $\Sigma_{i,j}$ is the covariance matrix of the reprojection error. In [1] the covariance matrix depends on the scale at which the keypoints are observed. In our case we choose to estimate it using the median absolute deviation (MAD) [22] that is a robust estimator of the standard deviation of the reprojection error. We perform the optimization using the Levenberg-Marquardt algorithm on matches found between the current and the previous frame. Then we refine this twist with an approach similar to [1] by projecting map points, transformed with the estimated twist, in the current frame to search for additional matches and obtain a more accurate estimation. The object pose in frame i is then updated as ${}^w \mathbf{T}_{o_i} = \exp(\mathbf{\Pi}^w \boldsymbol{\xi}_{o_i}) {}^w \mathbf{T}_{o_{i-1}}$. This tracking procedure is repeated for all objects.

F. Dynamic Bundle Adjustment

The goal of classical bundle adjustment is to refine the camera trajectory and 3D points position estimation. The dynamic bundle adjustment has multiple goals. First, the refinement of the dynamic objects trajectory and their 3D points position, jointly with the camera trajectory and 3D static points position. Second, it allows to link the object and the camera trajectory, indeed if the bundle adjustment did not take into account dynamic objects, only the camera pose would have an impact on the object pose, which would not improve it. By taking into account dynamic points whose position is estimated over time we can use them to refine the camera pose, similarly to static points but with less accuracy since object pose estimation is more noisy. Finally, it allows us to apply a soft constrain on twists within a temporal window. Doing so we obtain smoother trajectories and velocities that are more physically plausible.

Our bundle adjustment cost function can be written as follows:

$$E(\{{}^w \boldsymbol{\xi}_o, {}^c \mathbf{T}_w, {}^w \mathbf{X}, {}^o \mathbf{X}\}) = \sum_{i,j} e_{stat}^{i,j} + \sum_{i,j} e_{dyna}^{i,j} + \sum_i e_{const}^i \quad (19)$$

where $e_{stat}^{i,j}$ is the classical static reprojection error:

$$e_{stat}^{i,j} = \rho(\|{}^i \mathbf{x}_j - p({}^{c_i} \mathbf{T}_w, {}^w \mathbf{X}_j)\|_{\Sigma_{i,j}^{-1}})$$

$e_{dyna}^{i,j}$ is a dynamic reprojection error:

$$e_{dyna}^{i,j} = \rho(\|{}^i \mathbf{x}_j - p({}^c \mathbf{T}_w \exp(\mathbf{\Pi}^w \boldsymbol{\xi}_{o_i}) {}^w \mathbf{T}_{o_i}, {}^o \mathbf{X}_j)\|_{\Sigma_{i,j}^{-1}})$$

where $\Sigma_{i,j}^{-1}$ is estimated using the MAD as in equation (18). And e_{const}^i is a constant velocity model that penalizes twists variations by linking 3 consecutive poses:

$$e_{const}^i = \rho(\|\mathbf{\Pi}^w \tilde{\boldsymbol{\xi}}_{o_{i+1}} - \mathbf{\Pi}^w \tilde{\boldsymbol{\xi}}_{o_i}\|_w)$$

where \mathbf{W} is a diagonal weight matrix used to balance the errors, ${}^w \tilde{\boldsymbol{\xi}}_{o_{i+1}}$ is the twist linking the poses $\exp(\mathbf{\Pi}^w \boldsymbol{\xi}_{o_i}) {}^w \mathbf{T}_{o_i}$ and $\exp(\mathbf{\Pi}^w \boldsymbol{\xi}_{o_{i+1}}) {}^w \mathbf{T}_{o_{i+1}}$ and ${}^w \tilde{\boldsymbol{\xi}}_{o_i}$ is the twist linking the poses $\exp(\mathbf{\Pi}^w \boldsymbol{\xi}_{o_{i-1}}) {}^w \mathbf{T}_{o_{i-1}}$ and $\exp(\mathbf{\Pi}^w \boldsymbol{\xi}_{o_i}) {}^w \mathbf{T}_{o_i}$. Those twists are

computed using the logmap from $\text{SE}(3)$ to $\text{se}(3)$ defined in [20] and can be written for ${}^w\tilde{\xi}_{o_{i+1}}$ as:

$${}^w\tilde{\xi}_{o_{i+1}} = \log(\exp((\mathbf{\Pi}^w \xi_{o_{i+1}})^w \mathbf{T}_{o_{i+1}})(\exp(\mathbf{\Pi}^w \xi_{o_i})^w \mathbf{T}_{o_i})^{-1})$$

This equation moves each pose while respecting the mechanical joints constraints. Optimizing it can be cumbersome however the Schur trick can be applied as its Hessian is sparse [5]. These equations are classically optimized on a set of local keyframes that share visual information, but in our case, inspired by [8] we chose to have 2 sets of keyframes: temporal and spatial. All frames are converted to temporal keyframes to improve the tracking of fast moving objects and be able to track an object as soon as it enters the field of view of the camera. Keyframes stay in the temporal set for a fixed duration (in our case 5 seconds) they are then culled more severely than in ORB-SLAM2. This allows us to keep a reasonable number of keyframes. We chose to optimize camera poses on the set of temporal and local keyframes, while object poses are only optimized on the set of temporal keyframes and fixed in all other keyframes. Doing so, we apply our constant motion model only on the temporal window, allowing clusters to accelerate or decelerate.

G. Computing the cost functions jacobians

To optimize the cost functions (18) and (19) with a Levenberg-Marquardt optimizer we need to compute their jacobian. First we compute the Jacobian of the cost function used for object tracking $E(o_i \xi_w)$. Using the chain rule and getting inspiration from [20] it can be shown that:

$$\begin{aligned} J_E &= \frac{\partial E(o_i \xi_w)}{\partial o_i \xi_w} \\ &= \frac{\partial p({}^c \mathbf{X})}{\partial {}^c \mathbf{X}} ({}^c \mathbf{X}^\top \otimes \mathbf{I}_3)(\mathbf{I}_4 \otimes {}^c \mathbf{R}_w)(\mathbf{I}_3 \otimes {}^w \mathbf{T}_o) \partial_{\exp} \mathbf{\Pi} \end{aligned} \quad (20)$$

where \otimes is the Kronecker product, \mathbf{I}_N is an identity matrix of size N , ${}^c \mathbf{R}_w$ is the rotation matrix of ${}^c \mathbf{T}_w$ and:

$$\partial_{\exp} = \frac{\partial \exp(\xi)}{\partial \xi} = \begin{pmatrix} \mathbf{0}_{3 \times 3} & -[\mathbf{e}_1]_\times \\ \mathbf{0}_{3 \times 3} & -[\mathbf{e}_2]_\times \\ \mathbf{0}_{3 \times 3} & -[\mathbf{e}_3]_\times \\ \mathbf{I}_3 & \mathbf{0}_{3 \times 3} \end{pmatrix} \quad (21)$$

where $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ is the canonical base of \mathbb{R}^3 . Then, we compute the jacobian of e_{dyna} in equation (19), which is very similar to the previous jacobian. The derivatives of the function with respect to camera poses and points are the same as for classical bundle adjustment [26]. For the object poses we can compute:

$$\begin{aligned} J_{e_{dyna}} &= \frac{\partial e_{dyna}}{\partial o_i \xi_w} \\ &= \frac{\partial p({}^c \mathbf{X})}{\partial {}^c \mathbf{X}} ({}^c \mathbf{X}^\top \otimes \mathbf{I}_3)(\mathbf{I}_4 \otimes {}^c \mathbf{R}_w)(\mathbf{I}_3 \otimes {}^w \mathbf{T}_o) \partial_{\exp} \mathbf{\Pi} \end{aligned} \quad (22)$$

Finally we compute the jacobian of the constant velocity constrain with respect to each of the 3 twists involved in the constrain:

$$J_{e_{const}} = \begin{pmatrix} \frac{\partial e_{const}}{\partial {}^w \xi_{o_{i-1}}} & \frac{\partial e_{const}}{\partial {}^w \xi_{o_i}} & \frac{\partial e_{const}}{\partial {}^w \xi_{o_{i+1}}} \end{pmatrix} \quad (23)$$

we only show here the left part of the jacobian as the other parts are similar.

$$\frac{\partial e_{const}}{\partial {}^w \xi_{o_{i-1}}} = \frac{\partial \log(\mathbf{T})}{\partial \mathbf{T}} (\mathbf{I}_4 \otimes \mathbf{R}) \partial_{\exp} \mathbf{\Pi} \quad (24)$$

with $\mathbf{T} = \exp((\mathbf{\Pi}^w \xi_{o_i})^w \mathbf{T}_{o_i})(\exp(\mathbf{\Pi}^w \xi_{o_{i-1}})^w \mathbf{T}_{o_{i-1}})^{-1}$, \mathbf{R} is the rotation matrix of $\exp((\mathbf{\Pi}^w \xi_{o_i})^w \mathbf{T}_{o_i})(\exp(\mathbf{\Pi}^w \xi_{o_{i-1}})^w \mathbf{T}_{o_{i-1}})^{-1}$ and the derivative of the logmap is given by [20].

V. EXPERIMENTS

In this section we present the experiments we conducted to test our approach. We evaluate both the accuracy of the camera pose estimation and of the object pose estimation.

A. Experiments details

Datasets. We evaluate our approach on the KITTI [27] tracking dataset which consists in multiple sequences recorded from a camera mounted on a car. This dataset is particularly interesting for our approach as it contains the ground truth for both the camera pose and for some objects poses such as vehicles. It should be noted that the segmentation network can yield an "unknown" class which we consider to be static, as the dynamic classes in the KITTI dataset (cars, bikes and pedestrians) are correctly segmented by the network.

Metrics. The metrics for the evaluation of SLAM systems are usually the absolute translation error (ATE) [28] and the relative pose error (RPE) [29]. For each sequence we report the translation and rotation parts of the RPE, as it is done by both VDO-SLAM and DynaSLAM2. The object pose estimation accuracy can be evaluated using 2 different types of metrics: on the one hand the ATE and RPE that measure the quality of the objects trajectories and on the other hand the MOTP that evaluates the per-frame accuracy of objects 3D bounding boxes estimations and that we compute similarly to [5] using KITTI evaluation tools. As we do not estimate object boxes we use the ground truth box at the first pose of each object and propagate it using our camera and object pose estimations. We evaluate the true positive rate (TP) and the MOTP using the projected 3D bounding box (2D), in bird view (BV) and in 3D.

B. Camera pose estimation

In this subsection we evaluate the accuracy of our camera pose estimation. Similarly to [5] we only show here sequences in which the camera is moving. As we can see in table I our approach improves camera pose estimation on several sequences. To evaluate the stability of our approach we also computed the standard deviation over 10 runs for the sequence 3 and obtained a value of 1×10^{-4} and 4×10^{-4} for the translation and the rotation respectively. As objects are often either small, only visible for a short time or static, [1] performs well, but as we track clusters using many points, with a good precision, especially for clusters that do not move, we are able to reduce the drift. [6] also gives good results but requires depth information while our approach gives similar or better results than RGB based approaches. The most important improvement of our approach is in terms of object tracking accuracy as we can see in table II.

TABLE I
CAMERA POSE ESTIMATION COMPARISON ON THE KITTI TRACKING DATASET.

seq	ORB-SLAM2 [1]		DynaSLAM [3]		VDO-SLAM (RGB-D) [6]		DynaSLAM2 [5]		Ours	
	RPE _t (m/f)	RPE _R (°/f)	RPE _t (m/f)	RPE _R (°/f)	RPE _t (m/f)	RPE _R (°/f)	RPE _t (m/f)	RPE _R (°/f)	RPE _t (m/f)	RPE _R (°/f)
00	0.04	0.06	0.04	0.06	0.07	0.07	0.04	0.06	0.04	0.05
01	0.05	0.04	0.05	0.04	0.04	0.12	0.05	0.04	0.04	0.03
02	0.04	0.03	0.04	0.03	0.02	0.04	0.04	0.02	0.03	0.03
03	0.07	0.04	0.07	0.04	0.03	0.08	0.06	0.04	0.06	0.02
04	0.07	0.06	0.07	0.06	0.05	0.11	0.07	0.06	0.06	0.04
05	0.06	0.03	0.06	0.03	0.02	0.09	0.06	0.03	0.06	0.02
06	0.02	0.04	0.02	0.04	0.05	0.02	0.02	0.01	0.02	0.04
07	0.05	0.07	0.05	0.07	-	-	0.05	0.07	0.04	0.04
08	0.08	0.04	0.08	0.04	-	-	0.10	0.04	0.07	0.03
09	0.06	0.05	0.06	0.05	-	-	0.06	0.06	0.05	0.04
10	0.07	0.04	0.07	0.04	-	-	0.07	0.03	0.07	0.03
11	0.04	0.03	0.04	0.03	-	-	0.04	0.03	0.03	0.02
13	0.04	0.05	0.04	0.05	-	-	0.04	0.04	0.03	0.04
14	0.03	0.08	0.03	0.08	-	-	0.03	0.08	0.03	0.06
18	0.05	0.03	0.05	0.03	0.02	0.07	0.05	0.02	0.04	0.02
19	0.05	0.03	0.05	0.03	-	-	0.05	0.02	0.03	0.03
20	0.11	0.07	0.05	0.04	0.03	0.17	0.07	0.04	0.04	0.03
mean	0.055	0.046	0.051	0.045	-	-	0.053	0.041	0.044	0.034
std	0.020	0.016	0.016	0.015	-	-	0.019	0.020	0.015	0.011

C. Object pose estimation.

In this subsection we evaluate the accuracy of our object pose estimation. As we can see we improve object tracking accuracy, particularly for static objects such as the car 35 from sequence 11. The most important improvements usually come from the rotational part of the RPE, which is understandable as we only have 1 degree of freedom for the rotation of cars. We also observe that the most challenging cases happen when an object starts and stays far from the camera (e.g. seq. 05 and 10) because object tracking uses 3D points triangulated from stereo matches that are imprecise when points are far from the camera. We argue that the bruteforce keypoint matching of [5] help them when few frame to frame matches can be found, which can happen when the object is far from the camera. Furthermore we have not implemented a way to relocalize an object that has been lost for multiple frames. Thus on some sequences (e.g. car 0 of seq. 11 and car 12 of seq. 20) in which the objects are alternatively far and close from the camera, we are only able to track them on a small portion of their trajectory. However, we can see that we are generally able to accurately track objects for most of their trajectory. During our experiments we saw that we were able to track pedestrians, despite the fact that they are not rigid. We believe that our approach works because pedestrians undergo small deformations around arms and legs. As for camera pose estimation, we computed the standard deviation of object pose estimation for the sequence 3 and obtained values of 3.2×10^{-2} , 2.8×10^{-2} and 4.7×10^{-2} for the APE, the translational RPE and the rotational RPE. We also evaluated our results with an ANOVA which shows a significant difference (p-values ≤ 0.1) for most of our experiments.

We also show some qualitative results for the mapping, the camera and object pose estimation. The results are visible in figure 4. We are able to track multiple objects on all sequences. The estimated speed is very close from the ground truth with a maximum difference of about 3 km/h which occurs when the

object is far from the camera or when it is created. Looking at the bounding boxes we can see that they coincide and thus that the poses are well estimated for near and far objects. As we can see in the middle figure we can accurately track non rigid objects such as the cyclist, as long as most of their surface is rigid. We also show on the right an example of both a tracked static car and a car that slows down and speeds up. As we can see the estimated speed is close to the 0 for the static car, making the dynamic keypoints act like static ones.

VI. CONCLUSION

In this paper we proposed a new stereo semantic dynamic SLAM system called TwistSLAM, able to estimate both the pose of the camera as well as to track all dynamic objects in the scene. Using mechanical joints between clusters we can constrain objects movements to physically possible movements, which allows us to improve both camera and objects pose estimation compared to the state of the art.

REFERENCES

- [1] R. Mur-Artal and J. D. Tardós, "ORB-SLAM2: An open-source SLAM system for monocular, stereo, and RGB-D cameras," *IEEE Trans. on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [2] J. Engel, T. Schöps, and D. Cremers, "LSD-SLAM: Large-scale direct monocular SLAM," in *European Conf. on computer vision*, 2014.
- [3] B. Bescos, J. M. Fácil, J. Civera, and J. Neira, "DynaSLAM: Tracking, mapping, and inpainting in dynamic scenes," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4076–4083, 2018.
- [4] 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 *IEEE/RSJ IROS*, 2018, pp. 1168–1174.
- [5] B. Bescos, C. Campos, J. D. Tardós, and J. Neira, "DynaSLAM II: Tightly-coupled multi-object tracking and SLAM," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5191–5198, 2021.
- [6] J. Zhang, M. Henein, R. Mahony, and V. Ila, "VDO-SLAM: a visual dynamic object-aware SLAM system," *arXiv:2005.11052*, 2020.
- [7] M. Runz, M. Buffier, and L. Agapito, "Maskfusion: Real-time recognition, tracking and reconstruction of multiple moving objects," in *IEEE Int Symp. on Mixed and Augmented Reality (ISMAR)*, 2018, pp. 10–20.

TABLE II
OBJECT POSE ESTIMATION COMPARISON ON THE KITTI TRACKING DATASET. ATE IS IN M, RPE_t IN M/M, RPE_R IN °/M, TP AND MOTP IN %.

seq / obj. id / class	DynaSLAM 2 [5]									TwistSLAM								
	ATE	RPE_t	RPE_R	2D TP	2D MOTP	BV TP	BV MOTP	3D TP	3D MOTP	ATE	RPE_t	RPE_R	2D TP	2D MOTP	BV TP	BV MOTP	3D TP	3D MOTP
03 / 1 / car	0.69	0.34	1.84	50.0	71.79	39.34	56.61	38.53	48.20	0.31	0.10	0.28	58.02	60.00	58.02	60.00	58.02	60.00
05 / 31 / car	0.51	0.26	13.5	28.96	60.30	14.48	46.84	11.45	34.20	0.35	0.19	0.58	30.84	35.00	30.84	35.00	30.84	35.00
10 / 0 / car	0.95	0.40	2.84	81.63	73.51	70.41	47.60	68.37	40.28	0.77	0.21	1.98	7.20	3.70	6.10	3.10	5.80	2.80
11 / 0 / car	1.05	0.43	12.51	72.65	74.78	61.66	50.74	52.28	47.35	0.17	0.23	0.23	29.61	32.50	29.61	32.50	29.61	32.50
11 / 35 / car	1.25	0.89	16.64	53.17	65.25	19.05	31.95	6.35	26.02	0.10	0.03	0.11	65.00	67.50	65.00	67.50	65.00	67.50
18 / 2 / car	1.10	0.30	9.27	86.36	74.81	67.05	45.47	62.12	34.80	0.21	0.27	0.66	84.67	87.50	84.67	87.50	84.67	87.50
18 / 3 / car	1.13	0.55	20.05	53.33	70.94	21.75	41.45	16.84	35.80	0.15	0.21	0.56	28.19	30.00	28.19	30.00	28.19	30.00
19 / 63 / car	0.86	1.45	48.80	35.26	63.50	29.48	45.69	26.48	33.89	0.28	2.17	1.08	65.93	70.00	65.93	70.00	36.26	20.64
19 / 72 / car	0.99	1.12	3.36	29.11	62.59	29.43	55.48	29.43	39.81	0.16	0.05	0.34	16.92	20.00	16.92	20.00	16.92	20.00
20 / 0 / car	0.56	0.45	1.30	63.68	78.54	43.78	45.00	31.84	46.15	0.17	0.20	0.72	84.75	87.50	84.75	87.50	84.75	87.50
20 / 12 / car	1.18	0.40	6.19	42.77	76.77	37.64	49.29	36.23	40.81	0.24	0.20	1.54	14.24	17.5	13.91	17.45	13.04	17.25
20 / 122 / car	0.87	0.72	5.75	34.90	78.76	34.51	48.05	29.02	44.43	0.17	0.02	0.07	84.94	87.50	84.94	87.50	84.94	87.50
mean	0.93	0.61	11.83	55.15	70.96	39.05	47.01	34.08	39.31	0.26	0.32	0.68	45.53	49.89	47.41	49.84	44.84	43.18
std	0.23	0.35	12.58	17.65	6.19	17.82	6.13	18.27	6.36	0.18	0.58	0.34	29.49	30.24	29.66	30.32	29.33	30.02

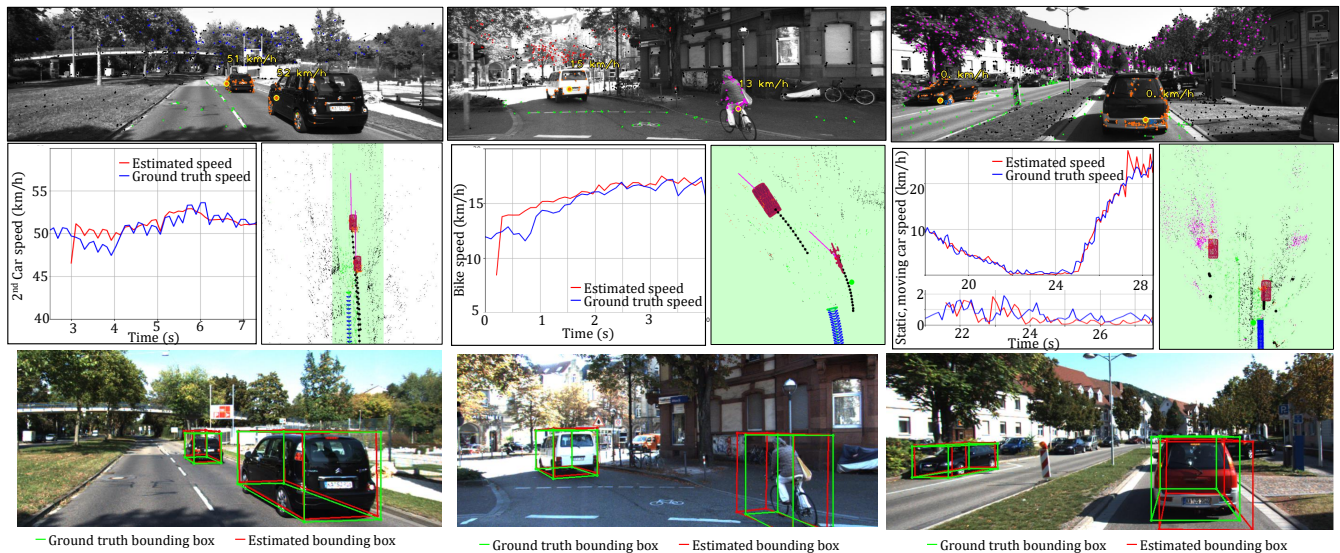


Fig. 4. Qualitative examples of cluster tracking on sequences from the KITTI tracking dataset. (Top) Frame with tracked clusters and their speed. (Middle left) Comparison of estimated (red) and ground truth (blue) speed (in km/h). (Middle right) Map with tracked clusters and camera poses, seen from above. (Bottom) Visualization of estimated poses (red) and ground truth (green) represented by their 3D bounding boxes.

- J. Huang, S. Yang, T.-J. Mu, and S.-M. Hu, "ClusterVO: Clustering moving instances and estimating visual odometry for self and surroundings," in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2020, pp. 2168–2177.
- C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Trans. on robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an open-source library for real-time metric-semantic localization and mapping," in *IEEE Int. Conf. on Robotics and Automation*, 2020, pp. 1689–1696.
- M. Gonzalez, E. Marchand, A. Kacete, and J. Royan, "S³LAM: Structured scene SLAM," *arXiv preprint arXiv:2109.07339*, 2021.
- S. Li and D. Lee, "RGB-D SLAM in dynamic environments using static point weighting," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2263–2270, 2017.
- S. Yang and S. Scherer, "CubeSLAM: Monocular 3-D object SLAM," *IEEE Trans. on Robotics*, vol. 35, no. 4, pp. 925–938, 2019.
- T. Taketomi, H. Uchiyama, and S. Ikeda, "Visual SLAM algorithms: A survey from 2010 to 2016," *IPSI Transactions on Computer Vision and Applications*, vol. 9, no. 1, pp. 1–11, 2017.
- E. Mouragnon, M. Lhuillier, M. Dhome, F. Dekeyser, and P. Sayd, "Real time localization and 3D reconstruction," in *IEEE Conf. on Computer Vision and Pattern Recognition*, vol. 1, 2006, pp. 363–370.
- G. Klein and D. Murray, "Parallel tracking and mapping for small AR workspaces," in *2007 6th IEEE and ACM int. symposium on mixed and augmented reality*, 2007, pp. 225–234.
- B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, "Bundle adjustment—a modern synthesis," in *Int. workshop on vision algorithms*. Springer, 1999, pp. 298–372.
- M. Rünz and L. Agapito, "Co-fusion: Real-time segmentation, tracking and fusion of multiple objects," in *2017 IEEE Int Conf. on Robotics and Automation (ICRA)*, 2017, pp. 4471–4478.
- J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *IEEE Conf. on computer vision and pattern recognition*, 2016, pp. 779–788.
- J.-L. Blanco, "A tutorial on se (3) transformation parameterizations and on-manifold optimization," *University of Malaga, Tech. Rep.*, vol. 3, p. 6, 2010.
- Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick, "Detectron2," <https://github.com/facebookresearch/detectron2>, 2019.
- E. Malis and E. Marchand, "Experiments with robust estimation techniques in real-time robot vision," in *2006 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2006, pp. 223–228.
- Z. Teed and J. Deng, "Raft: Recurrent all-pairs field transforms for optical flow," in *Eur. Conf. on Computer Vision*, 2020, pp. 402–419.
- A. I. Comport, É. Marchand, and F. Chaumette, "Kinematic sets for real-time robust articulated object tracking," *Image and Vision Computing*, vol. 25, no. 3, pp. 374–391, 2007.
- S. Axler, *Linear algebra done right*. Springer, 2014.
- F. Dellaert, "Visual slam tutorial: Bundle adjustment," *CVPR tutorial*, 2014.
- A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *2012 IEEE Conf. on computer vision and pattern recognition*, 2012, pp. 3354–3361.
- J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of RGB-D SLAM systems," in *Proc. of the Int. Conf. on Intelligent Robot Systems (IROS)*, Oct. 2012.
- Z. Zhang and D. Scaramuzza, "A tutorial on quantitative trajectory evaluation for visual (-inertial) odometry," in *2018 IEEE/RSJ Int Conf. on Intelligent Robots and Systems (IROS)*, 2018, pp. 7244–7251.