

A Framework for Depth Estimation and Relative Localization of Ground Robots using Computer Vision*

Rômulo T. Rodrigues¹, Pedro Miraldo², Dimos V. Dimarogonas², and A. Pedro Aguiar¹

Abstract—The 3D depth estimation and relative pose estimation problem within a decentralized architecture is a challenging problem that arises in missions that require coordination among multiple vision-controlled robots. The depth estimation problem aims at recovering the 3D information of the environment. The relative localization problem consists of estimating the relative pose between two robots, by sensing each other's pose or sharing information about the perceived environment. Most solutions for these problems use a set of discrete data without taking into account the chronological order of the events. This paper builds on recent results on continuous estimation to propose a framework that estimates the depth and relative pose between two non-holonomic vehicles. The basic idea consists in estimating the depth of the points by explicitly considering the dynamics of the camera mounted on a ground robot, and feeding the estimates of 3D points observed by both cameras in a filter that computes the relative pose between the robots. We evaluate the convergence for a set of simulated scenarios and show experimental results validating the proposed framework.

I. INTRODUCTION

The decentralized coordinated control of multiple mobile robots has several applications, including exploration [1], coverage [2], sampling and patrolling [3], search & rescue [4], and manipulation [5]. Achieving coordination within a decentralized architecture requires each agent in the formation to have some knowledge of the pose of its neighboring agents. The relative pose information allows to locally optimize a goal function, re-plan trajectories, and avoid collisions. However, for some mission scenarios such as mine caves or catastrophic areas, it is particularly challenging to obtain an accurate estimate of the relative pose due to the constrained communication channel bandwidth and impossibility to install or employ a high-precision positioning system, e.g., geolocation devices (GPS and Galileo) or Motion Capture systems. This paper aims at studying the problem of relative localization for two autonomous non-holonomic ground robots using 3D points observed by an

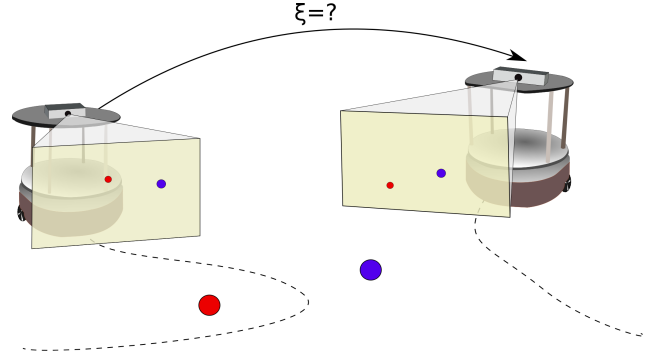


Fig. 1: Representation of the problem addressed in this paper. Two robots are observing a pair of 3D points. We propose a pipeline to simultaneously estimate the 3D depth of the two 3D points, while at the same time, obtain their relative pose.

onboard visual system (See Fig. 1). Since cameras observe 3D points up to an unknown scale, the formulation considers the problem of concurrently estimating the unknown depth of the 3D points.

In monocular computer vision, the mapping task aims at estimating the inherent 3D structure of the world from the 2D pixels in the image frame [6]. Considering two or more sequences of images taking from different views in a static environment, the depth of the points is computed by estimating the 3D points where the correspondent projection rays intersect. More sophisticated solutions have been proposed [7]–[9]. However, these works do not take into account the motion imparted to the camera. Instead it is considered a sequence of images captured at a known pose in the environment. Some authors study the estimation of depth maps from a single image (e.g. [10]–[12]). However, these works assume some high level knowledge of the environment. Recently, continuous 3D depth estimation has gained considerable attention (e.g. [13]–[17]). The goal is to minimize the estimation error of the 3D features over a continuous sequence of images ordered in time, by accounting for the dynamics of the camera. In contrast to other solutions that aim at algebraically estimating the unknown parameters as a fitting problem, continuous estimation considers the error dynamics and its convergence to the origin over time. This strategy benefits from: 1) being able to provide a metric about the quality of their estimates, e.g. by measuring the error's covariance or convergence rate; and 2) having robust performance by continuously improving the estimation based on new views. These advantages are particularly important

*This work was supported in part by SYSTEC UID/EEA/00147/2019 funded by national funds through the FCT/MCTES (PIDDAC); Project STRIDE NORTE-01-0145-FEDER-000033, supported by NORTE 2020, under the Portugal 2020 Partnership Agreement, through the European Regional Development Fund (ERDF); PDMA - NORTE-08-5369-FSE-000061, through NORTE 2020, IMPROVE POCI-01-0145-FEDER-031823, funded by FEDER funds through COMPETE2020, POCI and PIDDAC; COIN from Swedish Foundation for Strategic Research, Knut och Alice Wallenberg Foundation (KAW); and by the EU H2020 Co4Robots project.

¹R. T. Rodrigues A. P. Aguiar are with the Research Center for Systems and Technologies (SYSTEC), Faculty of Engineering, University of Porto, Porto, Portugal.

E-Mail: rtr@fc.up.pt and pedro.aguiar@fe.up.pt.

²P. Miraldo and D. V. Dimarogonas are with the Division of Decision and Control Systems, KTH Royal Institute of Technology, Stockholm, Sweden. E-Mail: {miraldo, dimos}@kth.se.

for image based visual servoing schemes [18], [19], in which the depth of the 3D features are fed in the control loop.

Relative pose estimation is another active topic in the robotics and computer vision communities. From [20], geometrically one can see that the minimum required number of 2D point correspondences is five. The number of required point correspondences is reduced even more if the motion is constrained in a plane. In this case, only two 2D point correspondences are required [21]. One of the disadvantages of these methods is that they estimate the translation parameters up to a scale. This means that one needs to get additional information about the observed environment to get the complete relative pose. Over the years, there have been many approaches to solve the relative pose problem, such as [22]–[24]. However, all of these methods were built on the assumption that the data (2D images) were discretely acquired over time, without any continuously time assumption. In [25], the authors propose a continuous perception strategy to estimate the relative localization of a set of robots in formation. Their method is based on the bearing-rigidity of the formation, and requires that the at least a pair robots should be in mutual visibility. In our work, we also aim at estimating the relative pose of the robots as in [25].

The scenario considered in this paper includes two mobile robots, navigating autonomously in an unknown environment. Both robots are equipped with perspective cameras. While previous solutions for this problem consider a set of two or more images from the environment or use some special fleet configuration (e.g. the robots are in each others' fields of view or have the ability of sensing the bearing-information about each others' positions), in this paper we propose a framework that shares a set of common observations of the environment in the respective local frame of each robot (3D point features are employed). A graphical representation of the problem is shown in Fig. 1. In this scenario, as will be shown in the experimental results, the proposed solution is able to estimate the 3 degrees of freedom (DoF) relative pose between the two robots using 3D points which are being continuously estimated by the two robotic agents.

To compute the relative pose, we use the fact that the two agents are tracking in the image plane two point correspondences. The proposed pipeline estimates the depth of the points seen by both robots while computing their relative pose. In contrast to [21], by taking account a continuous estimation of the relative transformation, our method is able to get the complete relative pose (the translation is not estimated up to a scale). As the robots navigate in the environment, they will be estimating the depth of the points being tracked in their respective image frames using a depth estimation filter. Then, an Extended Kalman Filter (EKF) receives the estimated depths to estimate the relative pose of the robots.

This paper is organized as follows. The next section presents an overview of the system proposed in this paper. In Sec. III we describe the proposed method for depth and relative pose estimation. Results with simulated and real data

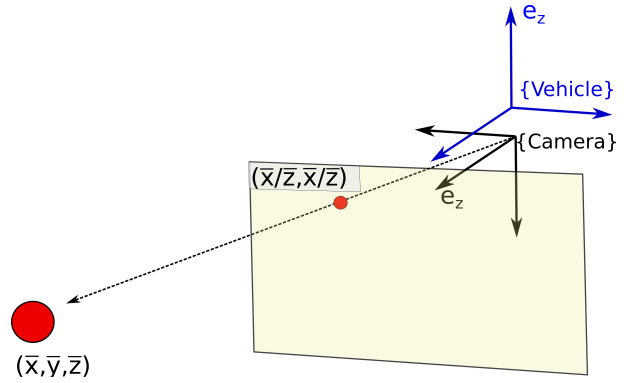


Fig. 2: Configuration of the body-fixed and camera frames, and the notations used for the projection model.

are shown in Sec. IV, and in Sec. V we discuss and conclude the paper.

II. SYSTEM OVERVIEW

This section presents the problem studied in this paper, an overview of the proposed system architecture, and the basic mathematical notation that holds for the remainder.

The problem discussed here is two-fold: 1) each agent has to compute the unknown depth of the points to recover the 3D structure of the world; and 2) each agent has to estimate the relative pose of its neighboring agent. The robot and camera configuration is shown in Fig. 2, and the proposed pipeline solution is shown in Fig. 3.

The cameras are mounted at the front of a non-holonomic vehicle, in which the optical axis of the sensor is aligned with the forward moving direction (see Fig. 2). Without loss of generality, we assume that the camera lies in the origin of the vehicle. The robot actuation, denoted as $\mathbf{u} = [v_d, w_\theta]^T \in \mathbb{R}^2$, consists of a forward and an angular speed command. This translates into a linear speed along the optical axis of the camera and an angular speed around the y -axis of the camera.

A depth estimation filter runs in each agent independently, recovering the depth of the 3D points. The depth filter, discussed in Sec. III-A, is an particularization of [14], constrained for cameras mounted on non-holonomic planar vehicles. Then, an Extended Kalman Filter (EKF) gathers the motion of each camera and the coordinates of the observed points to estimate the relative pose between the vehicles. This requires that one vehicle sends its 3D points estimation and actuation commands to its neighboring vehicle through the network. This is discussed in Sec. III-B. The EKF allows to capture the uncertainty inherent to the problem: depths converge online and measurements and actuation commands are subjected to noise.

The proposed pipeline includes two additional modules: image processing and network. The former is related with the detection and tracking of a 3D target on the image plane. For that purpose, one can use techniques such as SIFT [26], SURF [27], or ORB [28]. As a proof of concept, a color scheme detection is employed. It also eases the match of the correspondences between the points detected

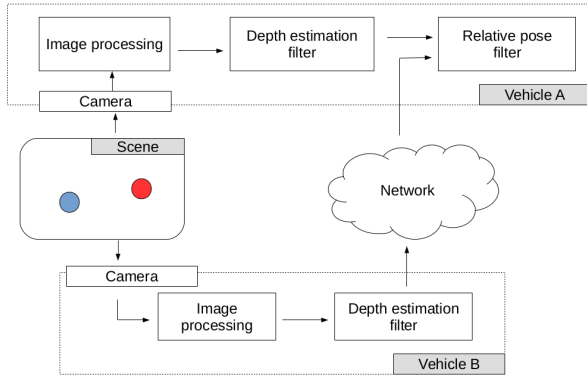


Fig. 3: Pipeline of the solution proposed in this paper. Each robot is equipped with an RGB camera that captures images of two 3D points. Then, each robot has a depth estimation filter that estimates the depth of the observed points. All the estimates are shared through the network (including the robot's control input), and a filter estimates the relative pose between the robots.

by the two robots. The network module enable the robots to communicate with each other. In practice, the vehicles are equipped with WiFi.

The mathematical notation adopted in this document is as follows: vectors are written in lower case bold font; matrices are written in upper case font; coordinate frames are denoted within curly brackets; and a subscript indicates the coordinate frame that a vector is described in (e.g. \mathbf{u}_A stands for the control input of camera A).

III. METHODS

This section presents the methods proposed for the simultaneous depth and relative pose localization of ground robots. We start by describing the used depth filter (Sec. III-A) for a single vehicle/point. Then, the relative pose filter is presented (Sec. III-B).

A. Depth Filter

This method is based on the framework proposed in [13], adapted to the motion of ground vehicles. Consider a camera moving freely in the 3D space, the dynamic of a point $\mathbf{p} = [\bar{x}, \bar{y}, \bar{z}]^T$ relative to the camera frame is

$$\dot{\mathbf{p}} = -\mathbf{w} \times \mathbf{p} - \mathbf{v} = \begin{cases} \dot{\bar{x}} = \bar{y}w_z - \bar{z}w_y - v_x \\ \dot{\bar{y}} = \bar{z}w_x - \bar{x}w_z - v_y \\ \dot{\bar{z}} = \bar{x}w_y - \bar{y}w_x - v_z \end{cases}$$

where $\mathbf{v} = [v_x, v_y, v_z]^T \in \mathbb{R}^3$ and $\mathbf{w} = [w_x, w_y, w_z]^T \in \mathbb{R}^3$ are the camera linear and angular velocity described in the camera frame. Since only the direction of the point is known, it is common practice to use the normalized image coordinates $\mathbf{s} = [\bar{x}/\bar{z}, \bar{y}/\bar{z}]^T = [x, y]^T$. The previous equation can be re-stated as

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{bmatrix} -\frac{1}{\bar{z}} & 0 & \frac{x}{\bar{z}} & xy & -(1+x^2) & y \\ 0 & -\frac{1}{\bar{z}} & \frac{y}{\bar{z}} & 1+y^2 & -xy & -x \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \mathbf{w} \end{bmatrix}.$$

As discussed in Sec. II, the vehicle input $\mathbf{u} = [v_d, w_\theta]^T$ consists in a forward speed (v_d) and an angular rate (w_θ). The forward direction of the vehicle is aligned with the optical axis of the camera, i.e., the linear velocity expressed in the camera frame satisfies $v_x = v_y = 0$ and $v_z = v_d$. The rotational axis of the vehicle is aligned with the y-axis of the camera such that $w_x = w_z = 0$ and $w_y = -w_\theta$. The reader is referred to Fig. 2 for a graphical illustration of the coordination frames at stake. Substituting $\chi = 1/\bar{z}$ to obtain a linear expression and simplifying the previous equation by removing the components multiplied by 0, the following dynamic system is obtained:

$$\begin{aligned} \dot{\mathbf{s}} &= f_m(\mathbf{s}, \mathbf{u}, t) + \Omega^T(t)\chi \\ \dot{\chi} &= f_u(\mathbf{s}, \mathbf{u}, \chi, t), \end{aligned}$$

where

$$\begin{cases} f_m(\mathbf{s}, \mathbf{u}, t) &= w_\theta \begin{bmatrix} (1+x^2) \\ xy \end{bmatrix} \\ \Omega(t) &= v_d \begin{bmatrix} x & y \end{bmatrix} \\ f_u(\mathbf{s}, \mathbf{u}, \chi, t) &= v_d\chi^2 + xw_\theta\chi \end{cases}. \quad (1)$$

In order to estimate the unknown depth \bar{z} , the estimation variables $\hat{\mathbf{s}}$ and $\hat{\chi}$ are introduced, as well as the associated estimation errors $\tilde{\mathbf{s}} = \mathbf{s} - \hat{\mathbf{s}}$ and $\tilde{\chi} = \chi - \hat{\chi}$. For that purpose, the following observer is considered:

$$\dot{\hat{\mathbf{s}}} = f_m(\mathbf{s}, \mathbf{u}, t) + \Omega^T(t)\hat{\chi} + H\tilde{\mathbf{s}} \quad (2)$$

$$\dot{\hat{\chi}} = f_u(\mathbf{s}, \mathbf{u}, \chi, t) + \Lambda\Omega^T(t)Q\tilde{\mathbf{s}}, \quad (3)$$

where $\Lambda \in \mathbb{R}_{++}$ is a positive gain and $H, Q = \alpha I_2 \in \mathbb{S}^2$ are symmetric positive matrices with $\alpha > 0$. Since for non-null v_d the camera moves along the optical axis, $\dot{\chi} \neq 0$. For this setup, the dynamics of the error can be proven to be locally exponentially stable. For a formal proof of the convergence of error to the origin, the reader is referred to [13]. The basic idea consists in guarantee that a Persistency of Excitation conditions holds. More precisely, this requires that the matrix $\Omega(t)\Omega(t)^T$ is full rank throughout time, that is

$$\int_t^{t+T} \Omega(\tau)\Omega(\tau)^T d\tau \geq \gamma I > 0, \quad t \geq 0,$$

where $\gamma > 0$ and I is an identity matrix of appropriate dimension. In particular, for a single point, this can be done by ensuring that $\|\Omega(t)\|^2 > 0$. From (1), we have

$$\|\Omega(t)\|^2 = (x^2 + y^2)v_d^2.$$

Therefore, the filter converges, as long as the projected point is not in the center of the image frame, i.e., $x \neq 0$ or $y \neq 0$ must hold for convergence. Geometrically, this constraint is easy to verify: if $x = y = 0$ and the camera moves along the z -axis, the motion will be along along the projection ray meaning and it would be impossible to recover the depth (not enough information to do triangulation).

In the next section, the recovered 3D points are employed for online relative pose estimation.

TABLE I: List of Notations for the relative pose estimation filter.

Variable	Description
A, B	Vehicles
$\{A\}, \{B\}$	Body-fixed coordinate frames
$\mathbf{u}_A, \mathbf{u}_B$	Control input of camera A and B
$(\mathbf{s}_i^A)_{i=1}^m, (\mathbf{s}_i^B)_{i=1}^m$	Normalized coordinates of the 3D points
$(Z_i^A)_{i=1}^m, (Z_i^B)_{i=1}^m$	Depth of the 3D points
ξ	Pose of B w.r.t A described in $\{A\}$

B. Relative Pose Filter

This section presents relative pose filter, which takes as input the depth estimation and the control inputs of the robots. First, consider that each robot has a body-fixed coordinate frame, denoted as $\{A\}$ for robot A and $\{B\}$ for robot B . The transformation from the body-fixed to the camera fixed frame is known for each robot. Now, instead of a single point, the robots observe m -points in the scenario. For each observed 3D point, each robot runs an instance of the depth estimation filter, as presented before. The pose of B with respect to A described in $\{A\}$ is denoted as $\xi = [\mathbf{r}^T, \theta]^T \in SE(2)$, where $\mathbf{r} = [r_x, r_y]^T$ is a 2D translation and θ an orientation angle. In accordance with the definition presented in Sec. II, the control input for each vehicle is defined as $\mathbf{u}_A = [v_{A,d}, w_{A,\theta}]^T$ and $\mathbf{u}_B = [v_{B,d}, w_{B,\theta}]^T$. The list of notations is summarized in Table I.

In order to derive the dynamic model of the relative pose between the robots, let $\xi_A = [\mathbf{r}_A^T, \theta_A]^T$ and $\xi_B = [\mathbf{r}_B^T, \theta_B]^T$ be the pose of robot A and B , respectively, described in an inertial coordinate frame. The dynamic of each robot is

$$\begin{cases} \dot{r}_{A,x} = v_{A,d} \cos \theta_A \\ \dot{r}_{A,y} = v_{A,d} \sin \theta_A \\ \dot{\theta}_A = w_{A,\theta} \end{cases} \quad \text{and} \quad \begin{cases} \dot{r}_{B,x} = v_{B,d} \cos \theta_B \\ \dot{r}_{B,y} = v_{B,d} \sin \theta_B \\ \dot{\theta}_B = w_{B,\theta} \end{cases}. \quad (4)$$

The relative pose between the two non-holonomic vehicles described in $\{A\}$ is given by

$$\xi = \begin{bmatrix} \mathbf{r} \\ \theta \end{bmatrix} = \begin{bmatrix} R^T(\theta_A)(\mathbf{r}_B - \mathbf{r}_A) \\ \theta_B - \theta_A \end{bmatrix}, \quad (5)$$

where the rotation matrix $R(\cdot)$ will be soon defined. The dynamics of the relative pose is

$$\dot{\xi} = \begin{bmatrix} \dot{R}^T(\theta_A)(\mathbf{r}_B - \mathbf{r}_A) + R^T(\theta_A)(\dot{\mathbf{r}}_B - \dot{\mathbf{r}}_A) \\ \dot{\theta}_B - \dot{\theta}_A \end{bmatrix}, \quad (6)$$

where $\dot{R}(a) = R(a)S(\dot{a})$,

$$R(a) = \begin{bmatrix} \cos a & -\sin a \\ \sin a & \cos a \end{bmatrix}, \quad \text{and} \quad S(\dot{a}) = \begin{bmatrix} 0 & -\dot{a} \\ \dot{a} & 0 \end{bmatrix}. \quad (7)$$

Substituting (5) in (6) yields

$$\dot{\xi} = \begin{bmatrix} -S(\dot{\theta}_A)\mathbf{r} + R^T(\theta_A)(\dot{\mathbf{r}}_B - \dot{\mathbf{r}}_A) \\ \dot{\theta}_B - \dot{\theta}_A \end{bmatrix}. \quad (8)$$

Finally, applying (4) and (7) in the previous equation leads to the dynamic model

$$\dot{\xi} = f(\xi, \mathbf{u}_A, \mathbf{u}_B) = \begin{bmatrix} v_{B,d} \cos \theta - v_{A,d} + w_{A,\theta} r_y \\ v_{B,d} \sin \theta - w_{A,\theta} r_x \\ w_{B,\theta} - w_{A,\theta} \end{bmatrix}. \quad (9)$$

The output model considers the relationship between the pose of the robots and the observed 3D points in the body-fixed frame. From the geometric perspective, the following equation holds for each of the m -points:

$$\mathbf{r} = \bar{z}_i^A \mathbf{s}_i^A - R(\theta) \bar{z}_i^B \mathbf{s}_i^B,$$

for which the corresponding time derivative is

$$\dot{\mathbf{r}} = \dot{\bar{z}}_i^A \mathbf{s}_i^A + \bar{z}_i^A \dot{\mathbf{s}}_i^A - R(\theta)[S(\dot{\theta})\bar{z}_i^B \mathbf{s}_i^B + \dot{\bar{z}}_i^B \mathbf{s}_i^B + \bar{z}_i^B \dot{\mathbf{s}}_i^B].$$

By replacing $\dot{\bar{z}} = -\dot{\chi}/\chi^2$, it is possible to describe the relationship between the dynamics of 3D points, stated in (2) and (3), and the dynamics of the relative pose:

$$\dot{\mathbf{r}} = -\frac{\dot{\chi}_i^A}{(\chi_i^A)^2} \mathbf{s}_i^A + \frac{\dot{\mathbf{s}}_i^A}{\chi_i^A} - R(\theta) \left[S(\dot{\theta}) \frac{\mathbf{s}_i^B}{\chi_i^B} - \frac{\dot{\chi}_i^B}{(\chi_i^B)^2} \mathbf{s}_i^B + \frac{\dot{\mathbf{s}}_i^B}{\chi_i^B} \right].$$

If vehicle A periodically estimates $(\mathbf{s}_i^A, \chi_i^A)$ and their dynamics, and receives the corresponding estimates from vehicle B , the following observation model holds:

$$h_i(\xi) = R^T(\theta)[\dot{\bar{z}}_i^A \mathbf{s}_i^A + \bar{z}_i^A \dot{\mathbf{s}}_i^A].$$

The relative pose filter consists in an Extended Kalman Filter that assumes that the input of the vehicles and the observations are perturbed by Gaussian noise with zero mean. First, in the prediction step, the prior is obtained according to (9)

$$\dot{\xi}^- = f(\dot{\xi}, \mathbf{u}^A, \mathbf{u}^B)$$

Then, for each pair of measurements, the update model is:

$$\dot{\xi} = \dot{\xi}^- + K[\pi_i(t) - h_i(\dot{\xi}^-)],$$

where K is the Kalman gain and

$$\pi_i(t) = S(\dot{\theta})\bar{z}_i^B \mathbf{s}_i^B + \dot{\bar{z}}_i^B \mathbf{s}_i^B + \bar{z}_i^B \dot{\mathbf{s}}_i^B$$

is computed from the measurements received from the neighboring vehicle. The Kalman gain and the covariance matrix are updated following the regular EKF framework.

With this we conclude the relative pose filter description. This setup allows us to obtain an estimate of the unknown state. We stress that this method builds on the assumption that at least two 3D points correspondences are required to align two coordinate systems [29]. Although we focus on the critical case with only two points, in general, the EKF framework benefits from multiple observations and increase robustness against occlusions.

IV. RESULTS

In this section, we discuss simulated (Sec. IV-A) and experimental results (IV-B). Simulations were performed using the Robot Operating System (ROS) and Gazebo/Turtlebot simulator at 20 [Hz]. Similarly, for the experimental results, ROS and two TurtleBots were employed. A footage of this experiment is attached.

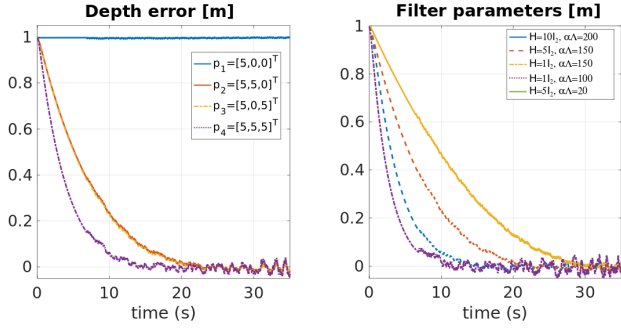


Fig. 4: Depth estimation for different scenarios using Gazebo/Turtlebot. Robot moves forward and does not rotate. At the left we show the depth estimation error for 4 different points. The graphic at the right shows the convergence behaviour for a single point and different filter parameters.

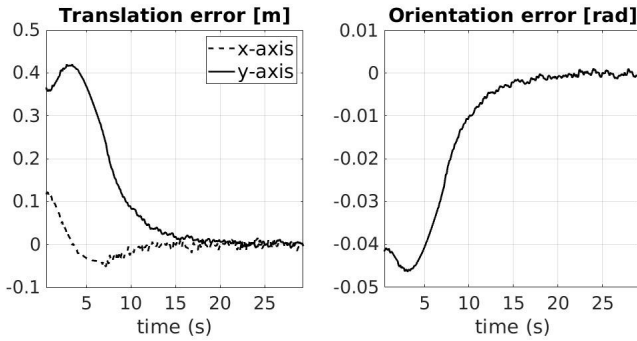


Fig. 5: Relative pose estimation: (a) translation and (b) orientation. The initial estimation error is 1.5 m in the x, y axes and 15 deg. in orientation.

A. Simulation

The first simulated scenario illustrates the depth filter properties presented in Sec. III-A. A single non-holonomic ground robot observes four 3D points. At $t = 0$, the points coordinates w.r.t. to the camera frame are $p_1 = [0, 0, 5]^T$, $p_2 = [0, 5, 5]^T$, $p_3 = [5, 0, 5]^T$, and $p_4 = [5, 5, 5]^T$. For each point, a filter was initialized with 1 [m] depth estimation error. The robot actuation is constant, set to $v_d = 0.1$ [m/s] and $w_\theta = 0$ [rad], that is, the robot moves straight ahead in the direction of the optical axis. Figure 4 shows the results for the three points using the same filter parameters. The estimation error converges to zero when the point does not rest at the origin of the image frame. In fact, the further the point is from the center of the image, the faster is the convergence due to the larger value of $\Omega(t)$. However, this also amplifies numerical errors and noise. Figure 4 at the right depicts the impact of different gains on the depth estimation, p_3 was selected. Larger values of H in (2) slow down the convergence rate, while higher values of $\alpha\Lambda$ increase the convergence rate. Here, once again, the designer must take into consideration that faster convergence leads to large steady-state errors.

Figure 5 shows the performance of the proposed depth estimation and relative pose estimation framework (method

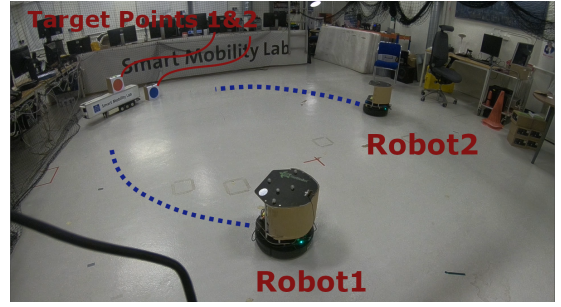


Fig. 6: Setup used in the experimental results. Two mobile robots are navigating in the environment while observing two 3D points (the red and blue targets).

proposed in Sec. III-B) altogether. Only the points p_1 and p_4 from the previous experience were considered and an additional robot was inserted in the scenario. For each set of robot/point, the initial depth estimation error was set to 1 [m] and the initial relative pose estimation error was set to 1.5 [m] in each transnational axis and 15° in orientation. The depth estimation error for each point follows a convergence behaviour similar to the one shown in Fig. 4. Therefore, as shown in Fig. 5, the translation estimation error converges to zero as the depth estimation error converges to zero. The same applies for the orientation error.

B. Experimental Results

Experiments were conducted at the Smart Mobility Lab¹, from KTH. We employed two TurtleBots (differential drive robots) named here *Robot_A* and *Robot_B*, respectively. Each robot was equipped with a perspective camera pointing towards the motion of the robot. The scenario contained two 3D targets (red and blue targets). For evaluation purpose, a motion capture system installed at the ceiling of the arena provides ground truth for the poses of the robots and the 3D points. Fig. 6 shows the proposed setup. The following velocities were sent to the robots: $v_d^A = 0.1$ [m/s] & $w_\theta^A = 0.02$ [rad/s], and $v_d^B = 0.1$ [m/s] & $w_\theta^B = -0.03$ [rad/s], for *Robot_A* and *Robot_B*, respectively. Through all the trajectory, we make ensure that the robots are observing the target points (around 40 [s]). The depth filter parameters were $H = 2.5I$ and $\alpha\Lambda = 120$ for both robots.

The results are shown in Fig. 7. The initial error for the depth filters were approximately 1 [m]. The depth of both points converge after 13 [s] for *Robot_A* and 25 [s] for *Robot_B*. The offset between the initial guess and the real values for the relative pose filter was 1.5 [m] in each of the translation axis and 60° in orientation. As one can see, after a transient period, the errors converge to zero, proving that the proposed pipeline is able to correctly estimate the relative pose between the two robots for the evaluated setup.

V. DISCUSSION

In this paper we presented a framework for depth and relative pose estimation for two vision-based robots. The

¹<https://www.kth.se/dcs/research/control-of-transport/smart-mobility-lab/smart-mobility-lab-1.441539> [accessed February 10, 2020].

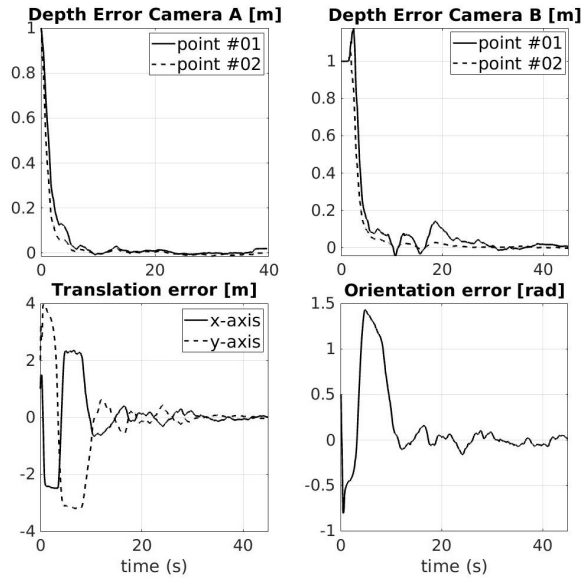


Fig. 7: Experimental results in a scenario with two 3D points: (a) and (b) shows the depth estimation error of the 3D points converging to the origin for robots *A* and *B*, respectively. Robot *A* estimates the pose of *B* described in $\{A\}$: (c) shows the relative translation estimation error and (d) the relative orientation estimation error.

depth estimation filter was built on top of recent results on continuous structure from motion methods. In particular, we constrained the problem to a camera mounted on a non-holonomic ground vehicle. The depth estimation and robot input commands are fed into an Extended Kalman Filter that estimates the relative pose between the ground robots. Simulations and experiments in a real scenario shown the convergence of both filters for several scenarios and set of initial conditions. In future work, the active control of the robots will be considered as well as the interplay between the gains and the position of the point in the image frame.

ACKNOWLEDGEMENT

The authors are grateful for the invaluable support provided by Pedro Roque (KTH) in the experimental setup.

REFERENCES

- [1] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A probabilistic approach to collaborative multi-robot localization," *Autonomous Robots*, vol. 8, no. 3, pp. 325–344, Jun 2000.
- [2] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Trans. Robotics and Automation (T-RA)*, vol. 20, no. 2, pp. 243–255, 2004.
- [3] A. Marino, G. Antonelli, A. P. Aguiar, A. Pascoal, and S. Chiaverini, "A decentralized strategy for multirobot sampling/patrolling: Theory and experiments," *IEEE Transactions on Control Systems Technology (T-CST)*, vol. 23, no. 1, pp. 313–322, 2015.
- [4] J. S. Jennings, G. Whelan, and W. F. Evans, "Cooperative search and rescue with a team of mobile robots," in *Int'l Conf. Advanced Robotics (ICAR)*, 1997, pp. 193–200.
- [5] D. Rus, B. Donald, and J. Jennings, "Moving furniture with teams of autonomous robots," in *IEEE/RSJ Int'l Conf. Intelligent Robots and Systems (IROS)*, vol. 1, 1995, pp. 235–242.
- [6] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, 2004.

- [7] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Trans. Robotics (T-RO)*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [8] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3D euclidean signed distance fields for on-board mav planning," in *IEEE/RSJ Int'l Conf. Intelligent Robots and Systems (IROS)*, 2017, pp. 1366–1373.
- [9] J. Engel, V. Koltun, and D. Cremers, "Direct sparse odometry," *IEEE Trans. Pattern Analysis and Machine Intelligence (T-PAMI)*, vol. 40, no. 3, pp. 611–625, 2018.
- [10] C. B. Choy, D. Xu, J. Gwak, K. Chen, and S. Savarese, "3D-R2N2: A unified approach for single and multi-view 3D object reconstruction," in *European Conf. Computer Vision (ECCV)*, 2016, pp. 628–644.
- [11] H. Fan, H. Su, and L. Guibas, "A point set generation network for 3D object reconstruction from a single image," in *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2463–2471.
- [12] K. Huang, Y. Wang, Z. Zhou, T. Ding, S. Gao, and Y. Ma, "Learning to parse wireframes in images of man-made environments," in *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 626–635.
- [13] R. Spica and P. R. Giordano, "A framework for active estimation: Application to structure from motion," in *IEEE Conf. Decision and Control (CDC)*, 2013, pp. 7647–7653.
- [14] R. Spica, P. R. Giordano, and F. Chaumette, "Active structure from motion: Application to point, sphere, and cylinder," *IEEE Trans. Robotics (T-RO)*, vol. 30, no. 6, pp. 1499–1513, Dec 2014.
- [15] R. Spica, P. R. Giordano, and F. Chaumette, "Plane estimation by active vision from point features and image moments," in *IEEE Int'l Conf. Robotics and Automation (ICRA)*, 2015, pp. 6003–6010.
- [16] A. Mateus, O. Tahri, and P. Miraldo, "Active structure-from-motion for 3d straight lines," in *IEEE/RSJ Int'l Conf. Intelligent Robots and Systems (IROS)*, 2018, pp. 5819–5825.
- [17] R. T. Rodrigues, M. Basiri, A. P. Aguiar, and P. Miraldo, "Low-level active visual navigation: Increasing robustness of vision-based localization using potential fields," *IEEE Robotics and Automation Letters (RA-L)*, vol. 3, no. 3, pp. 2079–2086, 2018.
- [18] F. Chaumette and S. Hutchinson, "Visual servo control. i. basic approaches," *IEEE Robotics Automation Magazine (RA-M)*, vol. 13, no. 4, pp. 82–90, 2006.
- [19] R. Spica, P. R. Giordano, and F. Chaumette, "Coupling active depth estimation and visual servoing via a large projection operator," *The International Journal of Robotics Research (IJRR)*, vol. 36, no. 11, pp. 1177–1194, 2017.
- [20] D. Nister, "An efficient solution to the five-point relative pose problem," *IEEE Trans. Pattern Analysis and Machine Intelligence (T-PAMI)*, vol. 26, no. 6, pp. 756–770, 2004.
- [21] D. Ortín and J. M. M. Montiel, "Indoor robot motion based on monocular images," *Robotica*, vol. 19, no. 3, p. 331342, 2001.
- [22] D. Scaramuzza and F. Fraundorfer, "Visual odometry [tutorial]," *IEEE Robotics Automation Magazine (RA-M)*, vol. 18, no. 4, pp. 80–92, 2011.
- [23] L. Kneip and H. Li, "Efficient computation of relative pose for multi-camera systems," in *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 446–453.
- [24] R. Ranftl and V. Koltun, "Deep fundamental matrix estimation," in *European Conf. Computer Vision (ECCV)*, 2018, pp. 292–309.
- [25] R. Spica and P. R. Giordano, "Active decentralized scale estimation for bearing-based localization," in *IEEE/RSJ Int'l Conf. Intelligent Robots and Systems (IROS)*, 2016, pp. 5084–5091.
- [26] D. G. Lowe, "Object recognition from local scale-invariant features," in *IEEE Int'l Conf. Computer Vision (ICCV)*, vol. 2, 1999, pp. 1150–1157.
- [27] H. Bay, A. Ess, T. Tuytelaars, and L. V. Gool, "Speeded-up robust features (SURF)," *Computer Vision and Image Understanding (CVIU)*, vol. 110, no. 3, pp. 346–359, 2008.
- [28] G. Bradski, K. Konolige, V. Rabaud, and E. Rublee, "ORB: An efficient alternative to SIFT or SURF," in *IEEE Int'l Conf. Computer Vision (ICCV)*, 2011, pp. 2564–2571.
- [29] P. Miraldo, S. Saha, and S. Ramalingam, "Minimal solvers for mini-loop closures in 3D multi-scan alignment," in *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2019.