

Panoptic-SLAM: Visual SLAM in Dynamic Environments using Panoptic Segmentation

Gabriel Fischer Abati^{1,2*}, João Carlos Virgolino Soares^{1*}, Vivian Suzano Medeiros^{1,3},
Marco Antonio Meggiolaro² and Claudio Semini¹

Abstract—The majority of visual SLAM systems are not robust in dynamic scenarios. The ones that deal with dynamic objects in the scenes usually rely on deep-learning-based methods to detect and filter these objects. However, these methods cannot deal with unknown moving objects. This work presents Panoptic-SLAM, an open-source visual SLAM system robust to dynamic environments, even in the presence of unknown objects. It uses panoptic segmentation to filter dynamic objects from the scene during the state estimation process. Panoptic-SLAM is based on ORB-SLAM3, a state-of-the-art SLAM system for static environments. The implementation was tested using real-world datasets and compared with several state-of-the-art systems from the literature, including DynaSLAM, DS-SLAM, SaD-SLAM, PVO and FusingPanoptic. For example, Panoptic-SLAM is on average four times more accurate than PVO, the most recent panoptic-based approach for visual SLAM. Also, experiments were performed using a quadruped robot with an RGB-D camera to test the applicability of our method in real-world scenarios. The tests were validated by a ground-truth created with a motion capture system.

I. INTRODUCTION

Simultaneous Localization and Mapping (SLAM) is a crucial part of the software frameworks used in many autonomous vehicles, such as drones and mobile robots, as it enables them to navigate and operate in previously unknown or unstructured environments. It consists of constructing a map of an unknown environment while simultaneously determining the position of the robot within that environment.

Visual SLAM systems are receiving increasing attention in the robotics community due to the low cost and richness of information provided by cameras. There are several visual SLAM systems in the literature with high precision using monocular (e.g. ORBSLAM [1], LSD-SLAM [2]), stereo (e.g. ORBSLAM2 [3]) and RGB-D cameras (e.g. RGBD-SLAM [4]).

However, these SLAM systems are not prepared to work in scenarios with moving/dynamic objects, which results in inaccurate localization and inconsistent mapping. To address this, there are several approaches to incorporate dynamic elements into visual SLAM. Recently, deep learning-based methods have been explored for SLAM in dynamic environments for providing high-level information about the scenes.

*The authors contributed equally to this work

¹ Dynamic Legged Systems Lab, Istituto Italiano di Tecnologia, Italy
gabriel.fischer@iit.it / joao.virgolino@iit.it / claudio.semini@iit.it

² Department of Mechanical Engineering at the Pontifical Catholic University of Rio de Janeiro, Brazil meggi@puc-rio.br

³ Department of Mechanical Engineering, University of São Paulo, Brazil viviansuzano@usp.br




Fig. 1: Overview of the Panoptic-SLAM for filtering dynamic objects. In (a) initial feature detection, (b) shows the output of panoptic segmentation, (c) shows the masks of moving instances (person and box), and (d) shows the keypoints belonging to moving objects successfully filtered, even though there is no cardboard box in the list of trained classes

Some methods, e.g. SGC-VSLAM [5], use object detection, such as YOLO [6] to detect and track labeled objects to then filter out the keypoints of the dynamic objects. DS-SLAM [7] and DCS-SLAM [8] use semantic segmentation combined with epipolar geometry to filter dynamic keypoints. Other methods, e.g. DynaSLAM [9], SAD-SLAM [10] and DOT-MASK [11] rely on instance segmentation techniques, such as Mask R-CNN [12] or YOLACT [13]. The main problem with all previous approaches is the necessity to have a pre-determined number of object classes that can be detected, and consequently, filtered. In other words, if a non-labeled object is moving in the scene, it would not be filtered and its features would cause a drift in the localization and outliers in the map.

Panoptic segmentation [14] is a computer vision task that combines both instance segmentation and semantic segmentation. In instance segmentation, objects in the image are identified individually and segmented at a pixel level, while in semantic segmentation, each pixel in the image is labeled with a semantic category. The goal of panoptic segmentation

is to unify these two tasks into a single one, where all pixels in the image are labeled with either an instance label, indicating which object it belongs to, or a semantic label, indicating which category it belongs to.

Recently some visual SLAM methods have been proposed using panoptic segmentation to handle dynamic scenarios. One example is PVO [15], which uses the visual odometry (VO) system from DROID-SLAM [16], and fuses panoptic segmentation with optical flow and depth measurements to create a panoptic-aware optical flow representation of the environment. Although its VO module can obtain accurate camera poses, the authors did not specify how the unknown objects from the panoptic segmentation model affects their system. Moreover, PVO can perform robustly in dynamic scenes, but it ignores loop closure detection when the camera returns to a previously visited location [15].

In this work, we present Panoptic-SLAM¹, a visual SLAM for dynamic environments that uses panoptic segmentation to detect and filter moving objects, as shown in Fig. 1. The contributions of the paper can be summarized as follows.

- A novel visual SLAM system based on panoptic segmentation that can perform robust localization in dynamic environments in the presence of unknown and unlabeled moving objects. The system can work with either monocular, stereo and RGB-D cameras since it only needs RGB image information.
- Tests using the Bonn RGB-D dynamic [17] dataset, which explicitly contains unknown dynamic objects in the scene, and with the TUM RGB-D dataset [18], obtaining the best results in several sequences compared to state-of-the-art methods, such as DynaSLAM [9], PVO [15] and FusingPanoptic [19].
- Experimental tests performed with an RGB-D camera attached to a quadruped robot moving in an indoor environment in a scenario containing moving and static unknown objects, such as other quadruped robots. A Vicon system was used to get the ground truth for accuracy evaluation. The new dataset was used to test the proposed approach, which proved to be on average 2 times more accurate compared to DynaSLAM.

II. RELATED WORK

Dealing with dynamic objects is a significant challenge in visual SLAM research. Most researchers treat dynamic objects as outliers, and several visual SLAM systems have been presented to address these outliers. These solutions can be broadly categorized into two main groups: geometry-based and learning-based methods. Geometry-based methods rely on classic computer vision techniques to detect and filter dynamic content. In general, they have inferior accuracy compared to learning-based methods. Their main advantage is to not require prior knowledge about the objects in the scene.

Learning-based methods, on the other hand, require a pre-trained model. They usually use object detection, seman-

tic segmentation, instance segmentation, or, more recently, panoptic segmentation. Li et al. [20] introduce a system that combines object detection and a geometry-based approach. The system is also based on ORB-SLAM2 with DeepSort tracking and epipolar geometry to detect static points of each object in the scene.

DS-SLAM [7] combines the SegNet semantic segmentation network with a moving consistency check algorithm based on optical flow to reduce the impact of dynamic objects, and generates a dense semantic octo-tree map from the environment.

Changing-SLAM [21] uses a YOLOv4 framework trained with the COCO dataset to detect objects in the scene and employs an extended Kalman filter to estimate their velocity, removing the keypoints belonging to moving objects. Although robust to both dynamic and changing environments, Changing-SLAM uses a feature repopulation method coupled with object detection instead of a segmentation network, and it does not present any specific mechanism to handle unknown objects.

None of the previously cited methods can deal with unknown moving objects. Ji et al. [22] presented a Semantic RGB-D SLAM approach that operates in dynamic environments by extracting semantic information solely from keyframes. Despite being able to deal with unknown objects using k-means and depth reprojection, their accuracy is lower than other methods, such as DynaSLAM, in environments with people. DynaSLAM [9] is one of the first works to use an instance segmentation model [12] to detect people in the scene at a pixel-level, which it then uses to filter dynamic features. It is one of the visual SLAM systems with the best accuracy in the TUM RGB-D dataset [18].

Yuan and Chen proposed SaD-SLAM [10], which integrates depth information and Mask R-CNN instance segmentation to identify dynamic features in images. The algorithm classifies each feature point as dynamic, static, or static and movable. SaD-SLAM exhibits high accuracy, surpassing DynaSLAM in certain scenarios. However, its main limitation is that the semantic segmentation is conducted offline. Vincent et al. presented DOT-Mask [11], which uses an extended Kalman filter to track dynamic objects in the scene, identified by an instance segmentation module. The authors aimed to develop a faster SLAM system at the expense of lower accuracy when compared to other approaches. The main problem of instance segmentation-based approaches for visual pose estimation is the lack of background information, which decreases the possibility to infer about the existence of unknown moving objects. The use of keypoints belonging to an unknown moving object for tracking might cause failures in pose estimation.

Recently, there has been an increasing use of panoptic segmentation in many applications. Semantic scene understanding with Lidar sensors [23] and 3D mapping [24] are examples of the usage of the panoptic segmentation outside the scope of SLAM. Panoptic segmentation is also used in visual state estimation systems, such as the previously mentioned PVO [15]. Analogous to PVO, SVG-LOOP [25]

¹<https://github.com/iit-DLSLab/Panoptic-SLAM>

presents a loop-closure detection algorithm that uses a combination of a semantic bag-of-words model processed with panoptic segmentation to minimize the impact of dynamic features, and a semantic landmark vector model to encode the geometric connections within the semantic graph.

Zhu et al. developed FusingPanoptic [19], a method based on ORB-SLAM2 that incorporates panoptic segmentation and geometry information. To minimize the impact of unknown moving objects, they propose a dynamic object classification approach based on epipolar geometry. Despite claiming robustness against unknown moving objects, this was not numerically evaluated in their paper. Also, they eliminate *a priori* all keypoints belonging to what they consider highly dynamic, without considering known movable objects that can be static (vehicles, for instance).

III. METHODOLOGY

Figure 2 shows a diagram of Panoptic-SLAM. The SLAM system is based on ORB-SLAM3 [26], and it is composed of four threads that run in parallel: panoptic segmentation, tracking, local mapping and loop closing. First, the image frames are processed in both tracking and panoptic segmentation threads. ORB features are extracted in the tracking thread, and the image is segmented into *Things*, *Stuff* and unknown objects in the panoptic segmentation thread. The *Things* objects are sent to a short-term data association algorithm to determine if they are new or were present in the last frame. The features associated with the background (*Stuff* objects) are matched with those of the last frame and used to compute a fundamental matrix. Using this fundamental matrix, the keypoints associated to known and unknown objects are classified as dynamic or static. Only static features are used for mapping and loop closing.

A. Panoptic Segmentation

In panoptic segmentation, pixels are classified either as *Things* or *Stuff*. *Things* are countable objects with well-defined boundaries and are potentially movable such as people, animals, or vehicles. *Stuff*, on the other hand, refers to uncountable amorphous regions of the image, mostly “unmovable” such as sky, floor, or walls [27]. The instance segmentation identifies individual objects in the *Things* category, while the semantic segmentation labels all pixels in the image with their corresponding category. With high accuracy and efficiency, this segmentation method surpasses previous box-based or box-free models.

Also, some parts of the image can result in being unknown, i.e. regions that the panoptic model could not assign any label to. The unknown masks can happen due to motion blur or unlabeled objects in the scene. There can also be wrong detections, i.e., objects with wrong classification. This can happen either if an object in the scene is labeled, but there is another class with similar characteristics (e.g. tv and monitor), or if an unlabeled object is similar to a labeled class.

Similar to [15] and [19], we have opted for PanopticFPN [28] as the panoptic segmentation model, although

any panoptic segmentation model can be employed in the Panoptic-SLAM framework. PanopticFPN is trained with the COCO dataset [29], which can segment up to 80 *Thing* different labels and 91 *Stuff* different labels. Figure 1(b) shows an example of an output of panoptic segmentation.

B. Dynamic Keypoint Filtering

The proposed method for dynamic keypoint detection and filtering is divided into four processes: *a priori* keypoint filtering, fundamental matrix computation, short-term data association, and moving keypoint classification of *Things* and Unknown.

Keypoint filtering: First, the panoptic model generates all masks predictions. People keypoints are filtered *a priori*, as humans can be considered as highly dynamic, and just in rare situations they remain completely still for a long period. The keypoints belonging to the *Stuff* mask are used for computing the fundamental matrix with RANSAC.

Fundamental matrix computation: In order to calculate the fundamental matrix, corresponding features must be identified in two consecutive frames. Feature matching involves finding pairs of points in the two frames that correspond to the same 3D point in the scene. Without an accurate feature matching, the fundamental matrix cannot be accurately calculated, which can result in errors in subsequent tasks. The feature matching algorithm matches the ORB descriptors of the keypoints in both the reference and the query images using a nearest neighbor approach. Specifically, for each keypoint in the reference image, the algorithm searches for the closest keypoint in the query image based on the similarity of their ORB descriptors.

The fundamental matrix is used to map feature points from the previous frame to their corresponding search domain in the current frame, namely the epipolar line. Assuming that the matched points in the current and last frames are p_1 and p_2 , respectively, their homogeneous coordinate form can be represented as $P1$ and $P2$.

$$P1_j = [u_1, v_1, 1], P2_j = [u_2, v_2, 1] \quad (1)$$

where u and v represent the pixel coordinates in the image frame, and j symbolizes the keypoint category (*Thing* or Unknown). Using these values, we can compute the epipolar line, denoted as L :

$$L \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = FP1_j \quad (2)$$

Given that X, Y, Z represent line vectors and F represents the fundamental matrix, the distance between a matched point and its corresponding epipolar line can be determined as

$$D_j = \frac{P2_j^T F P1_j}{\sqrt{\|X\|^2 + \|Y\|^2}} \quad (3)$$

where D_j is the distance between a matched point and its corresponding epipolar line. The threshold value for D_j was configured to 0.1, so the system could be sensitive to the movement of keypoints.

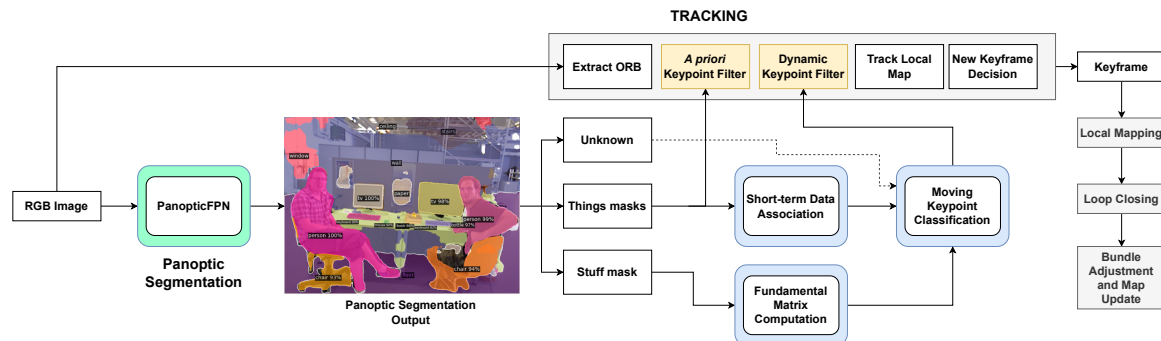


Fig. 2: Framework of Panoptic-SLAM. The processes highlighted in yellow are the additional modules included in ORB-SLAM3 to allow dynamic keypoint filtering. The processes highlighted in blue describe the dynamic keypoint classification, which also run in the tracking thread. To improve computational efficiency, panoptic segmentation (in green) runs in a separate thread

The idea is to have the fundamental matrix calculated with static keypoints and use it with *Things* and unknown feature matches in order to determine whether the distance from a matched point to its corresponding epipolar line is less than a certain threshold. If the distance is lower than the threshold, the matched keypoint is considered to be static.

Short-term data association: The moving keypoint classification of *Things* uses a short-term data association algorithm to address the issue of multiple objects with the same label in a frame. The short-term data association evaluates, using the intersection over union (IoU) metric, if a new Thing mask detected in the current frame correspond to a Thing mask of the last frame. Every instance mask is predicted with its respective bounding box. For every new pair of frames at time k and $k - 1$, the algorithm checks if any two bounding boxes from the same label overlap. Once an overlap is detected, the algorithm will extract the object contours C^k and C^{k-1} to compute the $IoU = |C^k \cap C^{k-1}| / |C^k \cup C^{k-1}|$ and determine if both contours have an association. To filter the keypoints of objects and determine which ones are in motion, it is essential to differentiate between multiple objects that share the same label. If an object in the current frame is not matched with any object in the previous frame, it is considered a new one and the features associated with its mask are filtered.

Moving keypoint classification: After the data association, the current and last frame keypoints belonging to *Things* from the same label and same tracking ID are matched. D_{thing} is calculated using the matched features and the fundamental matrix to determine which points are dynamic and, consequently, filtered.

Features belonging to unknown pixels are located by adding all known masks into a single image and analyzing the black areas in it, as shown in Fig. 3(a). The process is similar to the moving keypoint classification of *Things*. After matched keypoints from current and last frame are found in the unknown mask, these points are checked using the fundamental matrix to calculate the epipolar distance D_{unk} and to determine if they are static or dynamic. The unmatched points are filtered.

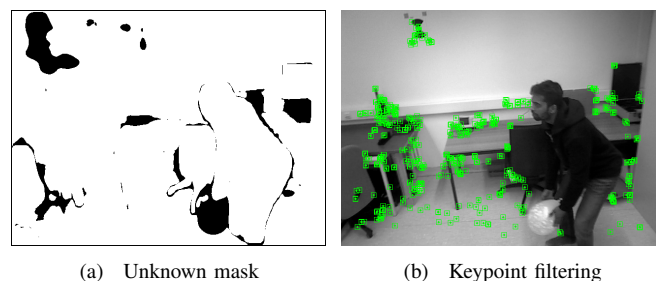


Fig. 3: Example of an unknown moving object being filtered in a sequence of the Bonn RGB-D dynamic dataset

IV. RESULTS

Sequences from two different benchmark datasets were used to evaluate the system and to compare it to state-of-the-art methods. Also, the system was tested in a real environment, measuring the ground-truth. The experimental results were also compared with other methods from the literature. All tests were performed five times and the median values were selected.

A. TUM RGB-D dataset

The TUM RGB-D dataset [18] is used to evaluate the robustness of the system in dynamic environments. It includes both RGB and depth image sequences captured by a Microsoft Kinect camera, along with corresponding ground truth trajectories. The data was recorded at a resolution of 640×480 and a frequency of 30 Hz. Four sequences were chosen for the evaluation: fr3_w_static, fr3_w_xyz, fr3_w_rpy and fr3_w_halfsphere. These sequences depict two people walking through a room, moving behind a desk, passing in front of the camera, and sitting on chairs. The main difference between each sequence is the camera motion. In the xyz sequence, the camera is moved along the three axis while maintaining a fixed orientation. In the rpy sequence the camera is rotated around roll, pitch, and yaw axis, maintaining a fixed position. In the halfsphere sequence, the camera follows a trajectory along a half-sphere. In

the static sequence, the camera remains still. The global consistency of the estimated trajectory is analyzed using the Absolute Trajectory Error (ATE) [18]. This metric compares the absolute distances between the translational components of ground truth and estimated trajectories.

Figures 4(a) and 4(b) show the comparison between the ground-truth and estimated trajectories of Panoptic-SLAM, respectively, in the `fr3_w_rpy` and `fr3_w_halfsphere` sequences. Despite being very challenging scenes, with fast and complex camera motions, especially in the halfsphere sequence, the Panoptic-SLAM system was able to perform an accurate estimation. Table I shows the ATE comparison between Panoptic-SLAM and several methods from the literature. The best results are highlighted in bold. Based on the comparison results, Panoptic-SLAM achieved a similar accuracy of DynaSLAM. This is likely due to the fact that DynaSLAM filters people in advance, whereas the TUM `fr3_walking` dataset primarily has people moving. Ji et al. [22] has an inferior accuracy than ours, despite also being robust to unknown labels in the environments. Our method also outperforms Zhu et al. [19], that also uses panoptic segmentation.

TABLE I: Comparison of the RMSE of ATE [m] of Panoptic-SLAM against 10 state-of-the-art methods using the TUM RGB-D dataset

Sequence	<code>fr3_w_static</code>	<code>fr3_w_xyz</code>	<code>fr3_w_rpy</code>	<code>fr3_w_half</code>
Panoptic-SLAM	0.009	0.014	0.032	0.025
ORB-SLAM3 [26]	0.038	0.819	0.957	0.315
PVO [15]	0.007	0.018	0.056	0.221
ReFusion [17]	0.017	0.099	—	0.104
DynaSLAM [9]	0.006	0.015	0.035	0.025
DS-SLAM [7]	0.008	0.024	0.444	0.030
SaD-SLAM [10]	0.017	0.017	0.032	0.026
DOTMask [11]	0.008	0.021	0.053	0.040
Ji et al. [22]	0.011	0.020	0.037	0.029
Zhu et al. [19]	0.013	0.018	0.039	0.030
DeFlowSLAM [30]	0.007	0.018	0.057	0.042

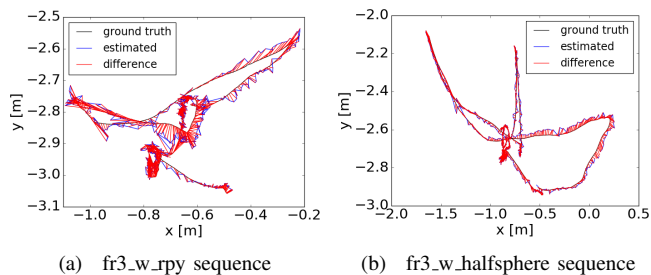


Fig. 4: Comparison between the ground-truth and the trajectory estimated by Panoptic-SLAM in two dynamic sequences of the TUM RGB-D dataset

B. Bonn RGB-D dynamic dataset

The Bonn RGB-D dynamic dataset [17] is used to evaluate the robustness against moving objects, people and non-labeled moving objects. It also uses the same evaluation

metrics from the TUM RGB-D dataset. Six sequences, filmed inside an indoor environment, were chosen from the dataset for evaluation: Balloon, Balloon2, non-obstructing box (non-obst box) 1 and 2, placing non-obstructing box (placing_no_box) 1 and 2. The Balloon sequences show a person walking and interacting with a floating balloon. In the non-obst box, a person moves a cardboard box from one place to another. In the placing_no_box sequences, a person appears and places a cardboard box on the floor. It is important to state that both balloon and cardboard box classes are not present in the COCO dataset [29] and, consequently, cannot be detected by our segmentation model.

Figure 5 shows the comparison between the ground-truth and estimated trajectories of ORB-SLAM3 and Panoptic-SLAM, respectively, in the moving non-obstructing box sequence. The trajectory estimated by ORB-SLAM3 completely deviates from the ground-truth, while Panoptic-SLAM is able to perform a correct estimation. Table II shows the RMSE of the ATE comparison between Panoptic-SLAM and ORB-SLAM3, DynaSLAM and ReFusion. The results of ReFusion and DynaSLAM were obtained in [17]. Our system outperformed ORB-SLAM3 in every sequence. Our system also outperformed the other systems in every sequence, except for the last one, where it achieved similar results to DynaSLAM, with a difference of 1 mm. Despite the fact that DynaSLAM achieved similar results to ours in the TUM RGB-D dataset, the same does not occur in the Bonn RGB-D dynamic dataset due to the presence of unknown moving objects. This is evident in the results of the non-obst box and placing-no-box sequences, where the error from DynaSLAM is 8.6 and 13 times larger than ours, respectively.

TABLE II: Comparison of the RMSE of ATE [m] of Panoptic-SLAM against ORB-SLAM3, DynaSLAM, and ReFusion using the Bonn RGB-D dynamic dataset

Sequence	Panoptic-SLAM	ORB-SLAM3	DynaSLAM	ReFusion
Non-obst box	0.027	0.347	0.232	0.071
Non-obst box2	0.033	0.043	0.039	0.179
Balloon	0.029	0.092	0.030	0.175
Balloon2	0.027	0.215	0.029	0.254
placing_no_box	0.044	0.842	0.575	0.106
placing_no_box2	0.022	0.023	0.021	0.141

To show the importance of each step of our methodology, experiments were made using three different configurations of the system: using only the people filter, the people filter together with the known moving object filter, and the people filter together with the unknown moving object filter. Two challenging sequences of the Bonn RGB-D dynamic dataset that contain unknown moving objects were chosen for the evaluation: non-obst box and placing-non-box. Table III shows the RMSE of the ATE of the different configurations, compared to the combined system. Due to the presence of both known and unknown moving objects in these scenes, every configuration fails except for the combined system. Using only the people filter with the unknown object filter, the

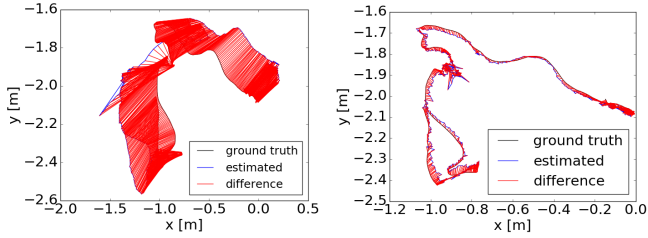


Fig. 5: Comparison between the ground-truth and the trajectory estimated by ORB-SLAM3 (left) and Panoptic-SLAM (right) in the Non-obstructing box sequence of the Bonn RGB-D dynamic dataset

results also drift because the cardboard box gets mislabeled in some frames, causing moving keypoints to not be filtered.

TABLE III: Evaluation of the ATE on the Bonn RGB-D dynamic dataset using Panoptic-SLAM with different configurations [m]

Sequence	non-obst box	placing_no_box
People filter	0.481	0.707
People + Moving object filter	0.029	0.765
People + Unknown object filter	0.302	0.721
Panoptic-SLAM	0.027	0.044

C. Experiments

Experiments were made to show the applicability of the proposed methodology in real-world indoor scenarios. The sequences were designed to represent real situations in a laboratory or a factory, with people and other robots moving, while the robot performing SLAM is also moving in the scene. Data was recorded in an Intel NUC computer attached to a Unitree AlienGo quadruped robot equipped with a RealSense D435 camera, as shown in Fig. 6 (left). A Vicon motion capture system was used to obtain the ground truth.

Three different sequences were recorded. In the first one, the robot moved in a static scenario, to serve as a baseline. In the second sequence, named `moving_robots`, the robot walked forward, while other robots walked nearby, namely a Spot from Boston Dynamics, and a Go1 from Unitree. Finally, in the third sequence, `robots_people`, the robot walked in a closed path, with not only robots walking near it, but also people, as shown by Fig. 6 (right).

Figures 7(a)-7(d) show the panoptic segmentation and keypoint filtering in two frames captured during the experiments. In Fig. 7(a) it is possible to differentiate the segmented people, the background and the quadruped robot Go1, which is unlabeled. With this information, it is possible to filter the keypoints belonging to the moving parts of the scene, as shown in Fig. 7(b). The same behavior is observable in Figs. 7(c) and 7(d).

Table IV shows the RMSE of the ATE comparison between Panoptic-SLAM, ORB-SLAM3, and DynaSLAM using the data recorded in the experiments. The best results are



Fig. 6: Images of the experiments: AlienGo quadruped robot with RGB-D camera (left). A scene of the `robots_people` sequence from an external point of view (right)

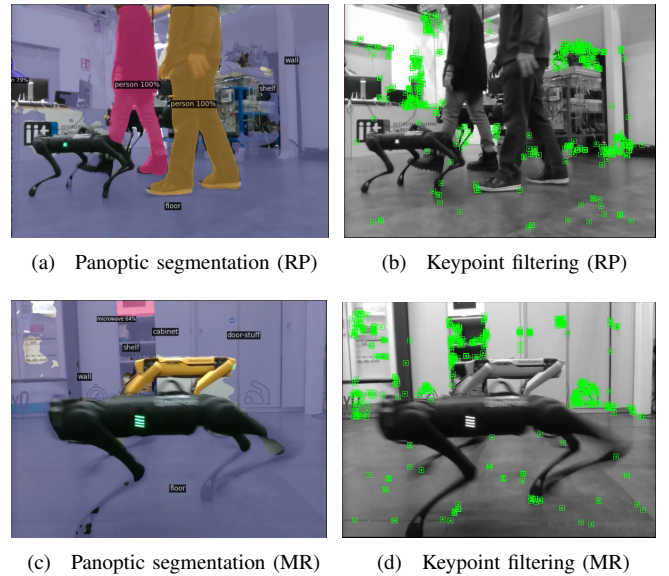


Fig. 7: Output examples of the panoptic segmentation and keypoint filtering steps from Panoptic-SLAM in the sequences `robots_people` (RP) and `moving_robots` (MR)

highlighted in bold. Our system outperformed ORB-SLAM3 in every sequence. As expected, the values for the baseline sequence are closer to each other. In the other sequences, on the other hand, ORB-SLAM3 has an increased error due to its incapability of filtering dynamic content. This is also evident in Figs. 8(a) - 8(f), that show the comparison between the ground-truth and estimated trajectories of ORB-SLAM3 and Panoptic-SLAM, respectively, in the experiments. Figures 8(c) and 8(e) show how the trajectories estimated by ORB-SLAM3 deviate from the ground-truth in several parts. Panoptic-SLAM is able to maintain a low-drift estimation, as shown in Figs. 8(d) and 8(f).

Figures 9(a) and 9(b) display the results from DynaSLAM in the `moving_robots` and `robots_people` experiments, respectively. Due to the lack of robustness to detect unlabeled moving objects, DynaSLAM could not accurately estimate its pose in the `moving_robots` sequence. In the other sequence, loop closure detection was important to lower the pose estimation drifts. Furthermore, DynaSLAM's people mask dilation helped to filter most of the dynamic content in

the scene. However, it also wrongly filters static keypoints. This can be seen by comparing the keypoint filtering from DynaSLAM and Panoptic-SLAM in Figs. 10(a) and 10(b). In the moving_robots sequence, on the other hand, DynaSLAM presented a higher error since it is not capable of filtering the dynamic keypoints due to the absence of people in the scene, as shown in Fig. 10(c).

TABLE IV: Comparison of the RMSE of ATE [m] of Panoptic-SLAM against ORB-SLAM3 and DynaSLAM in the experiments

Sequence	Panoptic-SLAM	ORB-SLAM3	DynaSLAM
baseline	0.057	0.062	0.064
moving_robots	0.090	0.495	0.291
robots_people	0.055	0.280	0.076

D. Run-time Analysis

All tests were performed on a notebook with Intel i7 CPU with 16 GB of RAM and NVIDIA RTX3060 GPU with 8 GB of VRAM running Ubuntu 20.04 LTS Linux. The SLAM system is implemented in C++, and the panoptic segmentation inference is implemented in Python, using the Detectron2 framework [31]. Table V shows the mean tracking time of the system in four sequences of the TUM RGB-D dataset, including the inference time of the panoptic segmentation. The average inference time of the panoptic segmentation was 0.2 second per frame for every sequence.

TABLE V: Mean tracking time [s] of Panoptic-SLAM in the sequences of the TUM RGB-D dataset

Sequence	Mean tracking time [s]
fr3_w_static	0.344
fr3_w_xyz	0.344
fr3_w_rpy	0.319
fr3_w_half	0.323

E. Limitations

As a limitation, the computational cost of the panoptic segmentation currently does not allow Panoptic-SLAM to run in real time. Furthermore, the panoptic segmentation results can have inconsistencies due to fast camera movements and changes in illumination. However, these inconsistencies have little effect on the overall result, since the filtering is based on the movement from the matched keypoints between frames, not the object classes.

V. CONCLUSIONS

This paper presented Panoptic-SLAM, an open-source visual SLAM system built on ORB-SLAM3 that operates online and is robust in dynamic environments, even in the presence of unknown and unlabeled moving objects. A panoptic scene segmentation is proposed in order to classify dynamic and static keypoints. We show the effectiveness of our method by comparing it with several systems from the

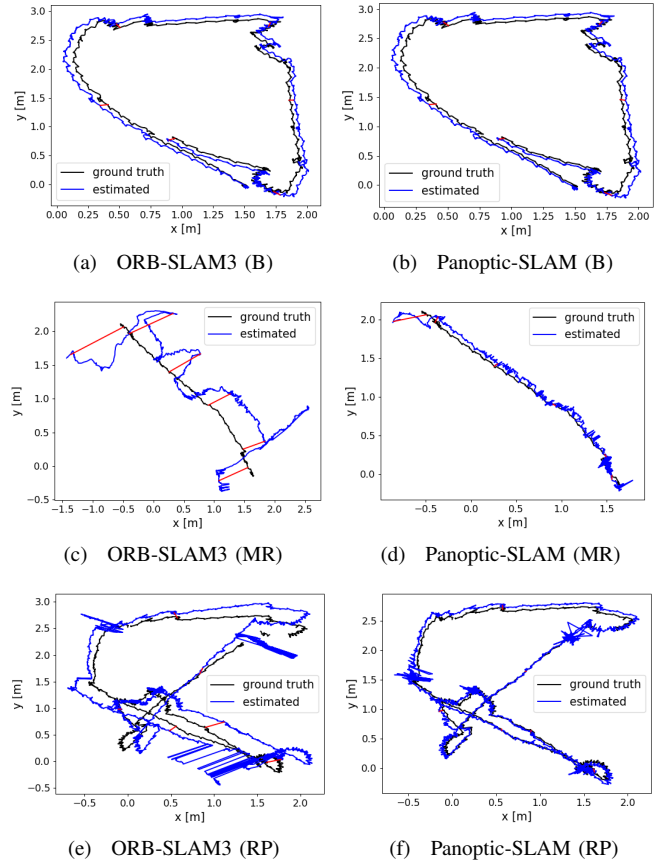


Fig. 8: Comparison between the ground-truth and the trajectory estimated by ORB-SLAM3 and Panoptic-SLAM in the sequences baseline (B), moving_robots (MR), and robots_people (RP)

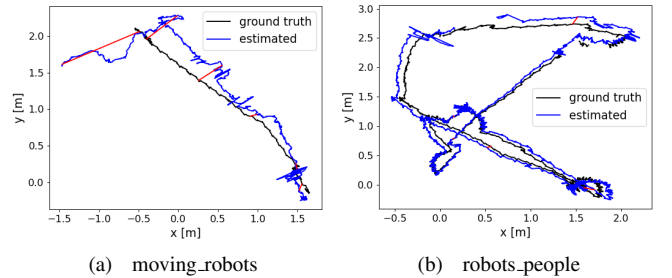


Fig. 9: Comparison between the ground-truth and the trajectory estimated by DynaSLAM in the moving_robots and robots_people sequences

literature that are considered to have the highest accuracy in dynamic environments, including DynaSLAM, DS-SLAM, SaD-SLAM, PVO and FusingPanoptic in challenging dynamic sequences. The results indicate that Panoptic-SLAM not only achieves the same levels of accuracy as DynaSLAM in highly dynamic scenes, but also surpasses it by a high margin in scenarios with unknown moving objects. Furthermore, Panoptic-SLAM was tested in experiments recorded with an RGB-D camera attached to a quadruped robot. The experiments consisted of highly dynamic scenes with people

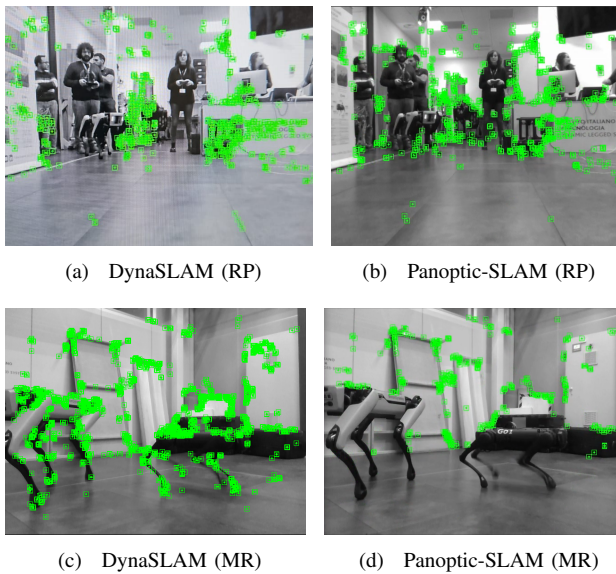


Fig. 10: Keypoint filtering comparison between DynaSLAM and Panoptic-SLAM in the sequences robots_people (RP) and moving_robots (MR)

moving with other quadruped robots. These robots were used to show the robustness of our system to filter non-labeled objects by the segmentation network.

For future works, we plan to deal with moving objects that occupy a large portion of the image in order to avoid lost tracks. Also, we aim to include in the method a semantic map building that can adapt over time, further exploring the capabilities of panoptic segmentation.

ACKNOWLEDGMENT

The authors acknowledge financial support from PNRR MUR Project PE000013 “Future Artificial Intelligence Research (FAIR)”, funded by the European Union – NextGenerationEU.

REFERENCES

- [1] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós, “ORB-SLAM: A Versatile and Accurate Monocular SLAM System,” *IEEE Transactions on Robotics*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [2] J. Engel, T. Schöps, and D. Cremers, “LSD-SLAM: Large-Scale Direct Monocular SLAM,” in *Computer Vision – ECCV 2014: 13th European Conference, Proceedings, Part II*, 2014, pp. 834–849.
- [3] R. Mur-Artal and J. D. Tardós, “ORB-SLAM2: An Open-Source SLAM System for Monocular, Stereo, and RGB-D Cameras,” *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [4] F. Endres, J. Hess, J. Sturm, D. Cremers, and W. Burgard, “3-D Mapping With an RGB-D Camera,” *IEEE Transactions on Robotics*, vol. 30, no. 1, pp. 177–187, 2014.
- [5] S. Yang, G. Fan, L. Bai, C. Zhao, and D. Li, “SGC-VSLAM: A Semantic and Geometric Constraints VSLAM for Dynamic Indoor Environments,” *Sensors*, vol. 20, no. 8, 2020.
- [6] J. Redmon and A. Farhadi, “Yolov3: An incremental improvement,” *arXiv preprint arXiv:1804.02767*, 2018.
- [7] 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 Int. Conf. on Int. Robots and Systems (IROS)*, 2018, pp. 1168–1174.
- [8] J. Zhang, M. Gao, Z. He, and Y. Yang, “DCS-SLAM: A Semantic SLAM with Moving Cluster towards Dynamic Environments,” in *IEEE Int. Conf. on Robotics and Biomimetics (ROBIO)*, 2022, pp. 1923–1928.
- [9] B. Bescos, J. M. Fàcil, 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.
- [10] X. Yuan and S. Chen, “SaD-SLAM: A Visual SLAM Based on Semantic and Depth Information,” in *IEEE/RSJ Int. Conf. on Int. Robots and Systems (IROS)*, 2020, pp. 4930–4935.
- [11] J. Vincent, M. Labb’e, J. Lauzon, F. Grondin, P. Comtois-Rivet, and F. Michaud, “Dynamic object tracking and masking for visual slam,” *IEEE/RSJ Int. Conf. on Int. Robots and Systems (IROS)*, pp. 4974–4979, 2020.
- [12] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask R-CNN,” in *Proceedings of the IEEE Int. Conf. on computer vision*, 2017, pp. 2961–2969.
- [13] D. Bolya, C. Zhou, F. Xiao, and Y. J. Lee, “Yolact: Real-time instance segmentation,” in *IEEE/CVF Int. Conf. on Computer Vision (ICCV)*, 2019, pp. 9156–9165.
- [14] A. Kirillov, K. He, R. Girshick, C. Rother, and P. Dollár, “Panoptic Segmentation,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 9396–9405.
- [15] W. Ye, X. Lan, S. Chen, Y. Ming, X. Yu, H. Bao, Z. Cui, and G. Zhang, “PVO: Panoptic Visual Odometry,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 9579–9589.
- [16] Z. Teed and J. Deng, “DROID-SLAM: Deep Visual SLAM for Monocular, Stereo, and RGB-D Cameras,” in *Neural Information Processing Systems*, 2021.
- [17] E. Palazzolo, J. Behley, P. Lotte, P. Giguère, and C. Stachniss, “ReFusion: 3D Reconstruction in Dynamic Environments for RGB-D Cameras Exploiting Residuals,” in *IEEE/RSJ Int. Conf. on Int. Robots and Systems (IROS)*, 2019, pp. 7855–7862.
- [18] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, “A benchmark for the evaluation of RGB-D SLAM systems,” in *IEEE/RSJ Int. Conf. on Int. Robots and Systems*, 2012, pp. 573–580.
- [19] H. Zhu, C. Yao, Z. Zhu, Z. Liu, and Z. Jia, “Fusing Panoptic Segmentation and Geometry Information for Robust Visual SLAM in Dynamic Environments,” in *IEEE 18th Int. Conf. on Automation Science and Engineering*, 2022.
- [20] G. Li and S. Chen, “Visual SLAM in Dynamic Scenes Based on Object Tracking and Static Points Detection,” *Journal of Intelligent & Robotic Systems*, vol. 104, no. 33, February 2022.
- [21] J. C. V. Soares, V. S. Medeiros, G. F. Abati, M. Becker, G. Caurin, M. Gattass, and M. A. Meggiolaro, “Visual localization and mapping in dynamic and changing environments,” *Journal of Intelligent & Robotic Systems*, vol. 109, 2023.
- [22] T. Ji, C. Wang, and L. Xie, “Towards Real-time Semantic RGB-D SLAM in Dynamic Environments,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2021, pp. 11175–11181.
- [23] B. Besic, N. Gosala, D. Cattaneo, and A. Valada, “Unsupervised Domain Adaptation for LiDAR Panoptic Segmentation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, p. 3404–3411, Apr. 2022.
- [24] Y. Miao, I. Armeni, M. Pollefeys, and D. Barath, “Volumetric Semantically Consistent 3D Panoptic Mapping,” *ArXiv*, vol. 2309.14737, 2023.
- [25] Z. Yuan, K. Xu, X. Zhou, B. Deng, and Y. Ma, “SVG-Loop: Semantic-Visual-Geometric Information-Based Loop Closure Detection,” *Remote. Sens.*, vol. 13, p. 3520, 2021.
- [26] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. M. Montiel, and J. D. Tardós, “ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual-Inertial, and Multimap SLAM,” *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [27] H. Caesar, J. Uijlings, and V. Ferrari, “Coco-stuff: Thing and stuff classes in context,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1209–1218.
- [28] A. Kirillov, R. Girshick, K. He, and P. Dollár, “Panoptic Feature Pyramid Networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6392–6401.
- [29] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, “Microsoft coco: Common objects in context,” in *European conference on computer vision*. Springer, 2014, pp. 740–755.
- [30] W. Ye, X. rong Yu, X. Lan, Y. Ming, J. Li, H. Bao, Z. Cui, and G. Zhang, “DeFlowSLAM: Self-Supervised Scene Motion Decomposition for Dynamic Dense SLAM,” *ArXiv*, vol. abs/2207.08794, 2022.
- [31] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick, “Detectron2,” <https://github.com/facebookresearch/detectron2>, 2019.