

# Robust Visual Teach and Repeat for UGVs Using 3D Semantic Maps

Mohammad Mahdavian, KangKang Yin and Mo Chen

**Abstract**—We propose a Visual Teach and Repeat (VTR) algorithm using semantic landmarks extracted from environmental objects for ground robots with fixed mount monocular cameras. The proposed algorithm is robust to changes in the starting pose of the camera/robot, where a pose is defined as the planar position plus the orientation around the vertical axis. VTR consists of a teach phase in which a robot moves in a prescribed path, and a repeat phase in which the robot tries to repeat the same path starting from the same or a different pose. Most available VTR algorithms are pose dependent and cannot perform well in the repeat phase when starting from an initial pose far from that of the teach phase. To achieve more robust pose independency, the key is to generate a 3D semantic map of the environment containing the camera trajectory and the positions of surrounding objects during the teach phase. For specific implementation, we use ORB-SLAM to collect the camera poses and the 3D point clouds of the environment, and YOLOv3 to detect objects in the environment. We then combine the two outputs to build the semantic map. In the repeat phase, we relocalize the robot based on the detected objects and the stored semantic map. The robot is then able to move toward the teach path, and repeat it in both forward and backward directions. We have tested the proposed algorithm in different scenarios and compared it with two most relevant recent studies. Also, we compared our algorithm with two image-based relocalization methods. One is purely based on ORB-SLAM and the other combines Superglue and RANSAC. The results show that our algorithm is much more robust with respect to pose variations as well as environmental alterations. Our code and data are available at the following Github page: [https://github.com/mmahdavian/semantic\\_visual\\_teach\\_repeat](https://github.com/mmahdavian/semantic_visual_teach_repeat).

**Index Terms**—Semantic Scene Understanding, SLAM, Vision-Based Navigation

## I. INTRODUCTION

**V**ISUAL Teach and Repeat (VTR) is an important task in robotic navigation. It has practical applications for repetitive tasks such as surveillance and transportation, especially for robots that are not equipped with GPS sensors, or in indoor areas where GPS sensors have low accuracy. VTR is thus an alternative for navigating an Unmanned Ground or Aerial Vehicle (UGV or UAV) at a relatively low cost, as it only requires a normal monocular camera.

VTR consists of a teach phase and a repeat phase. In the teach phase, the robot is driven by a user, or a path planner, and captures images along and around the path. During the repeat phase, the robot tries to repeat the same path starting from arbitrary locations and orientations, using only images captured by the camera. Figure. 1 demonstrates a teach path and multiple repeat paths using our robot in our testing area.

Manuscript received: February 24, 2022; Revised: May 20, 2022; Accepted: Jun 16, 2022.

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

The authors are with School of Computing Science, Simon Fraser University, Burnaby, Canada {mmahdavi, kkyin, mochen}@sfu.ca  
Digital Object Identifier (DOI): see top of this page



Fig. 1. An example of VTR paths starting from different locations and orientations.

There are two key challenges for VTR algorithms. First, during the repeat phase depending on the starting pose, the onboard camera may not be looking from the same location and angle as during the teach phase. In extreme cases such as transportation robots working in mines, the robot may need to move back to the starting point, so the camera position and orientation both change dramatically. Second, given that onboard memory is usually limited and robots cannot store all the raw images of the teach phase, VTR algorithms are desired to abstract the raw data to reduce storage as well as to relocalize in an online fashion.

The two common ways to process and abstract raw images for VTR are 1) creating 2D or 3D SLAM maps and 2) extracting and memorizing local features such as SIFT [1], SURF [2], BRIEF [3] or ORB [4] for each frame. The first way, requires advanced sensors such as Lidar, RGBD, or stereo cameras [10], which limits the application of the VTR. As an extreme version of VTR, we target UGVs that are only equipped with monocular cameras as a low-cost, light-weight, and ubiquitous sensor which can be mass-produced and quickly deployed. The SLAM maps generated using the monocular cameras are not highly accurate and in some cases scale ambiguous [6]. Therefore, we cannot solely rely on them to relocalize a camera in between any two runs. The second way, using local image features, only requires monocular cameras, but is problematic when the camera pose changes dramatically between the teach and repeat phases. Previous works using local features alone reported failures when viewing angles differ more than approximately 30 degrees [5]–[7].

**Contributions:** To the best of our knowledge, in the domain of VTR algorithms, we are the first to successfully use monocular cameras to obtain 3D semantic maps of the environmental objects to achieve more robust VTR. During the teach phase, we combine both SLAM and local image features to create 3D semantic maps of the environment. The maps contain both object locations and semantic labels, together with the 3D camera

poses on the teach path. In the repeat phase, our relocalization algorithm uses recognized objects as reference landmarks to align the camera poses to those in the teach map. Therefore the relocalization algorithm is robust with respect to starting camera pose changes, including planar position differences and the viewing angle variations, without requiring an accurate SLAM map. In addition, object movements between teach and repeat phases can be tolerated to a certain extent. Our results show that the repeat phase can start from a large range of initial poses in the environment, far away from the teach phase initial position with large viewing angle differences. This was not achievable by previous methods including using visual SLAM and image-based relocalization methods such as the combination of Superglue [26] and RANSAC [27]. We also demonstrate the ability to repeat the teach path in backward directions, which is not possible by any previous VTR work due to significant pose changes.

## II. RELATED WORKS

VTR has been investigated in a few studies for UAVs [8], [9] and UGVs [10], most of which are based on local image features. As one of the earliest studies, Furgale and Barfoot [11] built a manifold map of overlapping submaps as the robot was piloted along a route. The map was then used for localization as the robot repeats the route autonomously. Barfoot later developed a method called Multi-Experience Localization (MEL) by using local image features [12]–[17]. Such methods boost the performance by bridging between local image features found in multiple repeats. All these methods are sensitive to the initial viewing angle. The robot is possible to lose track of the path as the locations of local image features change drastically according to the viewpoint.

More recently, Camara et al. trained deep learning models to solve VTR with high end-point accuracy [18]. They used a CNN (Convolutional Neural Network) model as a feature extractor to build a feature database during the teach phase. At the repeat phase, the extracted features by the CNN model were compared with the feature database for visual place recognition. A horizontal offset estimator was also used to find the direction that the robot needs to move toward. However, the algorithm is still not robust to viewpoint changes, and may lose track of the path when starting from a location far away from the actual teach path.

Most recently, Dall'Osto et al. developed a VTR algorithm that utilized odometry information, paired with a correction signal driven by a computationally lightweight visual processing module [25]. Their algorithm cannot handle extreme initial pose changes, as will be shown in our comparative study.

In order to alleviate the pose-dependence problem, Ghasemi Toudeshki et al. proposed to utilize environmental objects [19], which can be robust semantic features of the environment that are viewpoint independent. Their algorithm memorized the semantic objects found during the teach phase for a UAV. Later a Seq-SLAM-like relocalization module used these objects to find the correct path toward the end-point. The work used objects as 2D image features without creating a 3D map, which resulted in inferior repeat accuracy compared to ours that uses environmental objects as 3D features. Again, we will show comparisons with this algorithm in our results section.

In order to solve the relocalization part of the VTR, there are other image-based algorithms available. For example, the

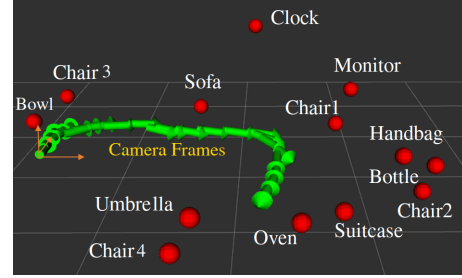


Fig. 2. An example of semantic map created in teach phase containing semantic object labels and positions as well as camera keyframes. Red dots represent the upper middle point of the corresponding objects.

combination of Superglue [26] and RANSAC [27] can be used for this purpose. Superglue is an image feature matching method based on a neural network that matches two sets of local features by jointly finding correspondences and rejecting non-matchable points. We combine it with RANSAC to find a transformation between teach and repeat phases initial poses. Also, ORB-SLAM [6] has relocalization capability. We will compare our algorithm with these two methods as well.

Our method shares the same motivation as [19] to use environmental objects as semantic features, but for the VTR problem of UGVs. Moreover, the 3D object locations are used to increase the accuracy of the algorithm, in addition to decreasing the sensitivity to initial robot poses. We were partly inspired by camera relocalization methods described in [5] and [24], which demonstrated that semantic objects can make camera location estimation more accurate and robust to significant viewing angle changes. As these methods were only applied to relocalization and not VTR, we did not perform a direct comparison with them. We note that our relocalization algorithm simplifies the one in [5] by using only 3D positions of the semantic objects. It is faster and more practical for VTR applications, and enables robust relocalization for even backward repeats not possible before. Although there are other possible methods for object matching like Hungarian algorithm [5] or RANSAC [27], we took a simple approach as it already works well for our case.

## III. PRELIMINARIES

### A. Overview of Proposed Method

We propose a VTR algorithm that is robust to large viewpoint changes, e.g., when the robot starts from a position far away from the teach path. The key of our method is to generate a semantic 3D map, rather than a complete SLAM occupancy grid map, from the environment containing observed objects as well as the camera trajectory. We use a SLAM module to track the camera path and a recognition module to detect objects from image frames. As an example implementation, we choose the well-known ORB-SLAM [6] to obtain camera poses from local image features, and a CNN-based model You Only Look Once (YOLOv3) [20] to recognize objects in the environment. Note that other choices may work within our framework as well. We combine the outputs of ORB-SLAM and YOLOv3 to build a 3D semantic map, as shown in figure 2, that contains semantic object labels, 3D object positions, along with the robot trajectory.

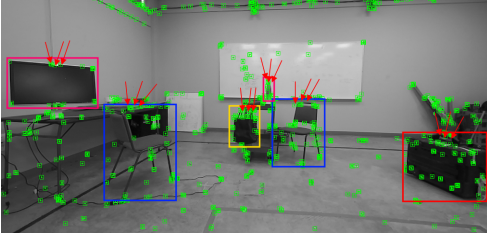


Fig. 3. Local features detected by ORB-SLAM and bounding boxes detected by YOLOv3. The closest image features within a threshold to the upper middle point of the bounding boxes, pointed by the red arrows, are used for calculating the 3D positions of the corresponding objects.

### B. ORB-SLAM

ORB-SLAM [6] has been widely adopted to reconstruct the camera trajectory and generate a sparse 3D reconstruction of the environment for monocular, stereo, and RGB-D cameras in real-time. It is also able to perform map reuse, loop closure, and relocalization with small viewing angle changes. In this work, we use ORB-SLAM for a monocular camera. The SLAM module gets initialized after receiving multiple image frames when the camera starts moving. The bundle adjustment algorithm takes in the motion of extracted 2D local ORB features and produces a map containing camera poses and 3D positions of these ORB features. Figure 3 shows an example of the 2D local ORB features detected during a robot motion. The origin of the generated map is the starting point where the ORB-SLAM module was initialized. However, the generated map is ambiguous in its scale. Even ORB-SLAM maps for the same area may have different scales in different runs. We will discuss in Section V-A how to handle this scale ambiguity in our relocalization algorithm.

We cannot solely rely on ORB-SLAM or any other similar visual SLAM algorithms to relocalize the camera in between any two different poses due to deficiency of similar local image features after dramatic view angle changes, as shown in Section VI-D. Therefore, we combine local features obtained from visual SLAM with an object detector to add the necessary semantic information for robust VTR.

### C. YOLOv3

To detect objects in each image frame, we employ a CNN-based model, YOLOv3 [20]. The YOLOv3 model applies a single neural network to the full image to extract image features first. Then it divides the image into several regions and predicts the location of the bounding box, the semantic label, and their confidence scores for each object. To achieve high accuracy, the network was trained on 24 most common objects in indoor areas, such as TV-Monitor, Sofa, Chair, Umbrella, Clock, Bottle, etc., from the COCO 2017 dataset. Also, the Darknet-ROS [21] package was utilized to publish objects detected by YOLOv3 to ROS. Figure 3 shows the detected objects by YOLOv3 during a robot motion.

YOLOv3 was chosen for its high accuracy as well as its fast inference speed, which is necessary since several modules need to run in parallel. In general, any other highly accurate and fast object detector could be used for our VTR algorithm.

## IV. TEACH PHASE

In the teach phase of VTR algorithms, the robot moves either with manual control from a user, or by a path planning algorithm. We call the trajectory of this motion the *teach path*, which the robot needs to memorize in order to repeat it later. To this end, we wish to create a 3D semantic map, referred to as the *teach map* hereafter, that contains camera keyframes, semantic object labels, and object positions in 3D. Camera keyframes can be obtained from ORB-SLAM directly, as described in Section III-B. Simultaneously, a YOLOv3 model is utilized to detect and classify objects in the streaming images to provide us the semantic object labels.

In order to obtain 3D object locations, we combine the 3D positions of ORB features from ORB-SLAM and the estimated object bounding boxes from YOLOv3. More specifically, we estimate the 3D object position from the 3D positions of the ORB features inside the corresponding object's bounding box. Empirically, the upper middle area of each object is the least occluded and most visible region of the object in image frames. Therefore, we first find the closest ORB features to the upper-middle point of the object bounding box within a threshold, and then use their average 3D positions as the object location in the semantic map. A similar heuristics was implemented in [5] to create a 3D semantic map of the environment to relocalize the camera poses. In figure 3, we use red arrows to point out the chosen ORB feature points for estimating 3D object locations.

Object locations estimated from different image frames are usually not identical numerically, even for the same object, due to noise and errors in camera movement and image feature detection. We therefore only add a semantic landmark and its object location into the semantic map when its estimated 3D position is beyond a reasonable threshold away from any previously added objects with a similar label. We also ignore bounding boxes that are only partially observable inside the image frames to avoid large estimation errors.

For simplicity, we ignore any duplicate objects during the teach phase and only keep track of unique objects in the semantic map. In principle, duplicate object classes can be considered, but could introduce unnecessary ambiguities in the semantic map. In practice, many unique objects exist in most indoor environments. In addition, as we demonstrate in Section VI, a very sparse semantic map containing 3D positions of just a few objects is sufficient for robust VTR.

By the end of the teach phase, we have obtained a semantic teach map containing 3D positions of environmental objects as well as their semantic labels, and a camera pose trajectory  $P^k$  where  $k \in \{1, \dots, K\}$  is the keyframe index. Figure 2 shows a sample semantic map together with the origin of the map for our testing scene used in the experiments.

## V. REPEAT PHASE

We aim to develop a repeat algorithm that is highly robust to large variations of the robot starting poses, such as faraway starting locations or viewing angles opposite to the ones of the teach phase. Tolerance to reasonable environment changes, such as relocation of a subset of objects, is desirable as well. The algorithm should also be able to repeat the teach path with reasonable accuracy, relative to the size of the robot.

The key to such a robust VTR algorithm is to accurately re-localize the robot pose with respect to the teach map coordinate



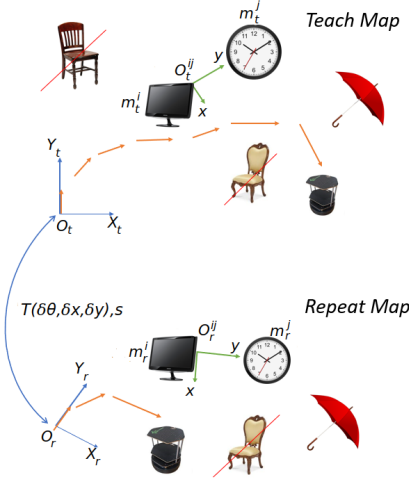


Fig. 4. Relocalization between the teach and repeat phase calculates  $\delta x, \delta y, \delta \theta$  and  $s$  from two matching objects.

frame. For this purpose, we developed an optimization-based relocalization algorithm that finds the **best matching pair of objects** between the teach and repeat maps to relocalize the robot robustly and accurately. After relocalization, the robot can simply move toward the closest point on the teach path and repeat it to the end or move back to the starting point.

#### A. Relocalization

To relocalize the robot in the repeat phase, 3D positions of objects in the environment estimated during the teach phase are used as reference landmarks to transform the robot pose in repeat map  $P_r$  to its counterpart in the teach map  $P_t$ . We note that the estimated 3D positions of the objects contain noise and errors from various sources, due to estimation errors from ORB-SLAM, occasional object relocation in the environment, and noise caused by the movement of the robot. Therefore, we robustly estimate the relative transformation from  $P_r$  to  $P_t$  using the best matching pair of objects in the environment. Furthermore, we assume there are multiple unique objects in the environment. We only consider unique objects to reduce ambiguity caused by repetitive objects, and leave it as future work to utilize repetitive objects for relocalization.

In general, there are 6 Degrees of Freedom (DoFs) for a given rigid body. However, our camera is mounted with a fixed base on a UGV that can only move in the horizontal ground plane and rotate around the vertical axis. Therefore, there are only 3 DoFs left for the robot/camera. In addition, ORB-SLAM outputs are ambiguous in scale. Thus there are in total four unknowns ( $s, \delta x, \delta y, \delta \theta$ ) in the relative transformation between two semantic maps, where  $s$  is the scalar,  $(\delta x, \delta y)$  are the planar translation, and  $\delta \theta$  is the rotation around the vertical axis. We thus require the positions of two matching environmental objects to solve for the four unknowns. The scalar  $s$  can be easily calculated by dividing the relative distances of the pair of objects in the two maps. We now detail how to solve for the relative transformation between  $P_r$  and  $P_t$  from intermediate coordinate frames estimated from the two matching objects.

Figure 4 illustrates various coordinate frames involved in relocalization, together with some environmental objects.  $O_t$  and  $O_r$  are the two 2D coordinate frames automatically

extracted by ORB-SLAM from the robot starting pose in teach and repeat phases, respectively. They define the map's origin and starting orientation. We first ignore repetitive objects in both maps, such as the chairs in figure 4. Then object pairs with the same labels in both maps are enumerated, e.g., the monitor and the clock. Hereafter we denote individual objects as  $m^i$  and  $m^j$ , and the corresponding object pair as  $m^{ij}$ . Next, a new 2D coordinate frame  $O^{ij}$  is defined from the chosen objects  $m^{ij}$  for both maps. The  $y$ -axis is the vector pointing from one object to the other in the horizontal plane, e.g., the vector from the monitor to the clock. The  $x$ -axis is the cross product of the  $y$ -axis and the up direction. We denote a 3-DoF transformation matrix  $T$  with one rotational DoF  $\alpha$  and two translational DoFs  $dx, dy$  as in (1). Then, the transformation of all objects from a map coordinate frame, either  $O_t$  or  $O_r$ , to its corresponding object coordinate frame, either  $O_t^{ij}$  or  $O_r^{ij}$ , is denoted by (2) and (3).

$$T(\alpha, dx, dy) = \begin{bmatrix} \cos \alpha & -\sin \alpha & 0 & dx \\ \sin \alpha & \cos \alpha & 0 & dy \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (1)$$

$$T(\alpha_t, dx_t, dy_t) = O_t^{ij} T \quad (2)$$

$$T(\alpha_r, dx_r, dy_r) = O_r^{ij} T \quad (3)$$

where the subscript  $t$  and  $r$  denote the various quantities in either the teach or repeat map. We denote the coordinates of all objects in a map as  $M_t = \{m_t^k, k \in (1 \cdots N_t)\}$  or  $M_r = \{m_r^k, k \in (1 \cdots N_r)\}$ . Then, using (4) and (5) we can transform an object  $m_t^k$  in  $M_t$  or  $m_r^k$  in  $M_r$ , in their original map coordinates to their corresponding object coordinates as follows:

$$O_t^{ij} m_t^k = O_t^{ij} T \times O_t m_t^k \quad (4)$$

$$O_r^{ij} m_r^k = O_r^{ij} T \times O_r m_r^k \quad (5)$$

Then the aggregated error  $\gamma$  between positions of all  $N$  unique objects that appear in both maps can be calculated:

$$\gamma(O_t^{ij}, M_t, O_r^{ij}, M_r, s^{ij}) = \sum_{k=1}^N \left( O_t^{ij} m_t^k - s^{ij} \times O_r^{ij} m_r^k \right)^2 \quad (6)$$

$s^{ij}$  is a scale factor introduced by ORB-SLAM, which can be estimated by:

$$s^{ij} = \left\| \overrightarrow{m_t^{ij}} \right\| / \left\| \overrightarrow{m_r^{ij}} \right\| \quad (7)$$

where  $\left\| \overrightarrow{m_t^{ij}} \right\|$  and  $\left\| \overrightarrow{m_r^{ij}} \right\|$  are the Euclidean distance between  $m^i$  and  $m^j$  in the teach and repeat maps, respectively.

The right-hand side of (6) would be zero for perfectly accurate maps without any object relocation. But in real-world applications it is nonzero, and we aim to find the best matching pair of objects, denoted as  $(m^{i*}, m^{j*})$ , that minimize this aggregated error  $\gamma$ . That is, we search for the object indices that minimize (6) as follows:

$$(i^*, j^*) = \arg \min_{(i,j) \in (1 \cdots N)} \gamma(\cdots) \quad (8)$$

In practice, to make our relocalization algorithm more robust to estimation errors and outliers caused by object movement,

we only use the top matching objects in calculating  $\gamma$ . Specifically, positional errors between two maps for all objects are first sorted in an ascending order. Only the top half of objects are used in (6) and minimizing (8).

After finding the best matching pair of objects between the two maps, the optimal scale ratio ( $s^*$ ) is calculated by (7). The optimal 3-DoF transformation matrix  $T^*$  between the teach and repeat map is calculated from the best matching pair of objects ( $m^{i^*}, m^{j^*}$ ) as follows:

$$O_r T^* (\delta\theta, \delta x, \delta y) = O_{i^* j^*} T \times O_r^{i^* j^*} T \quad (9)$$

Therefore, three of the unknowns ( $\delta x, \delta y$  and  $\delta\theta$ ) are solved for and embedded in the above matrix  $T^*$ . It is now straightforward to transform the robot pose in the repeat map  $P_r$  to the teach map pose  $P_t$  as follows:

$$P_t = s^* \times (O_r T^* \times P_r) \quad (10)$$

It is interesting to note that as long as fewer than half of the objects are moved, the object movements would be detected during the relocalization phase due to our previous assumption of only considering top matching objects.

Note that while introducing object-level data association may help the algorithm to work with repetitive objects, our tests showed that due to the highly different camera poses between teach and repeat phases, the objects' side that the camera is observing can be completely different and in many cases, there may be no similarity between two sides of an object. Although training a CNN-based model using the captured images of the teach phase, similar to [28], can help to better distinguish the objects during repeat phase, doing so is outside of the scope of this paper, as we do not save the captured images. In any case, one advantage of our algorithm is that no training is needed after the teach phase.

Another solution for this problem can be to memorize the relative topological position of the objects in the scene during teach phase and use them during the repeat phase. However, in some cases, moving and rotating similar objects in the environment may prevent the algorithm from distinguishing objects from object-level association or relative topological positioning, which is needed in our relocalization algorithm.

### B. Forward Motion

After transforming the initial robot pose from the repeat map to the teach map, the next step is to move toward the teach path and repeat it. The forward motion involves moving toward the closest point on the teach path to the robot, and then following the list of keyframes  $P^k$  captured in the teach phase toward the end. Therefore, the closest point on the teach path from the current pose of the robot is considered as the first goal point of the robot. Then the next goal points are chosen with an arbitrary distance to the current point within the range of one meter. The robot would first rotate toward the next goal point and then move from the current position toward the next one all the way to the endpoint  $P^K$ .

### C. Backward Motion

The robustness of our system is manifested by its ability to repeat the teach path in a backward direction. Our VTR algorithm relocalizes the robot from the repeat map to the teach map using detected semantic landmarks, independent of the viewpoint. Then instead of choosing the next points

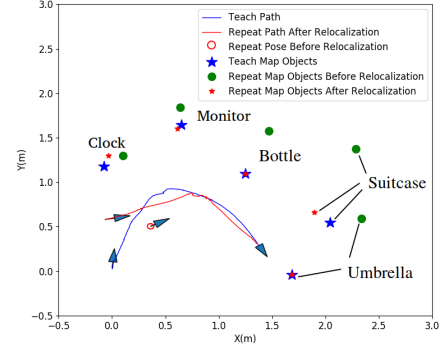


Fig. 5. An example of the teach map and the repeat map before and after relocalization. Here, the bottle and the umbrella were matched, and helped to relocalize the robot pose in repeat onto the teach map. The arrows show the initial robot pose.

toward the end-point on the teach path, the robot can choose points directed to the starting location on the teach map. Therefore, the backward motion is basically following the list of keyframes captured in the teach phase from the closest point on the teach path toward the starting keyframe  $P^1$ .

## VI. EXPERIMENTS

To evaluate the performance and accuracy of our algorithm and system, several experiments were carried out in a lab at Simon Fraser University. The lab was equipped with a Vicon motion capture system, which was only used for evaluation purposes and not used in our VTR method. The test platform was a Turtlebot2 [22] robot equipped with a ZED2 camera capturing only monocular data. The computer hardware was Intel CPU Core i9-9980HK and RTX 2080 GPU, running the ROS Kinetic on Ubuntu 16.04.

The robot was placed inside the experimental area in which a number of objects were randomly placed. Then the robot was controlled manually by a user along a teach path, while ORB-SLAM, YOLOv3, and the object position detection modules were running simultaneously. The VTR system also built the semantic map and memorized the teach path in an online fashion. An example semantic map built in the teach phase can be seen in figure 2.

During the repeat phase, we placed the robot in various initial positions and orientations in the lab. Then the robot moves slightly on a pre-defined short path, either autonomously, e.g., the initial curved motion in Repeat 3, 4 and 5 in figure 6 or manually, e.g., the initial part of the Repeat 1 and 8 in figure 6, to both start ORB-SLAM to track the camera pose and briefly observe the environment for the mapping modules to create a new semantic map for the repeat phase. The precise details of this initial motion is not too important: As long as a subset of at least three objects in the lab were seen, the relocalization module transforms the current robot pose to the teach map using the observed objects in both maps.

### A. Forward Repeat

Having relocalized the robot with respect to the 3D semantic map built in the teach phase and while the ORB-SLAM is running, the VTR system activated the motion planning module to move the robot toward the closest point on the teach

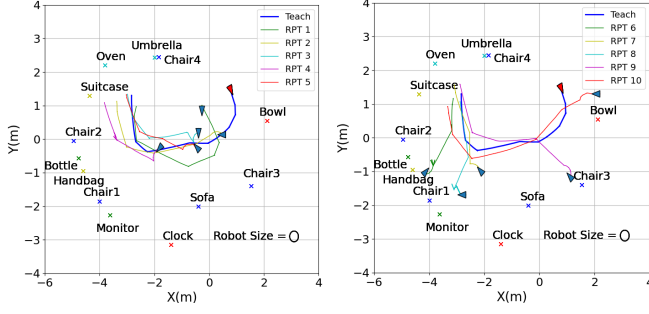


Fig. 6. The teach and ten repeat motions captured by our Vicon motion capture system for forward repeat tests. Triangle arrows indicate the initial poses of the robot, for which we chose to cover a large range of starting positions and orientations.

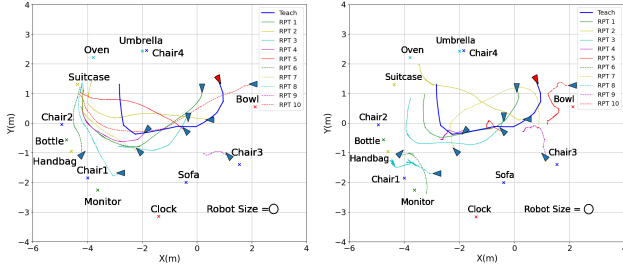


Fig. 7. The teach and repeat paths resulted from VTR algorithms published in [19](left) and [25](right), using similar initial robot poses as the tested poses for our own algorithm shown in figure 6

path and followed its keyframes toward the end. In figure 5 we show the object locations and robot paths: on the teach map; on the repeat map before relocalization; and on the repeat map after relocalization. In this specific test, the bottle and the umbrella were chosen by our algorithm to be the best matching object pair. The repeat path taken by the robot converged to the end-point with high accuracy, by traversing some of the camera keyframes of the teach path.

We have performed in total 10 forward repeat tests, all of which were captured by the Vicon motion-capture system and are shown in figure 6. The thick blue lines indicate the teach path, and all other colored lines visualize the repeat paths. The starting points for repeat motions were strategically chosen to cover a large range of locations and viewing angles inside our lab. In most cases, the viewing angles were significantly different from the one in the teach phase. The locations of the environmental objects utilized in the VTR algorithm are also shown in the figure. In all tests, the robot was first moved slightly to initialize the ORB-SLAM and semantic map for relocalization. Then the robot was able to repeat the teach path with reasonable accuracy, independent of the initial robot location and viewing angle. Also, to show our algorithm works in environments with few unique objects different tests have been performed and you can see its performance here: [https://www.youtube.com/watch?v=raRT7S9NSfc&list=PLuLzEWWNu1\\_qjPRa5OqTj6P4EjpY7veUA](https://www.youtube.com/watch?v=raRT7S9NSfc&list=PLuLzEWWNu1_qjPRa5OqTj6P4EjpY7veUA)

We report key statistics of the 10 repeat tests in Table I. The distances between the start point in the teach path and the ones in the repeat paths were calculated, and the average

TABLE I

Statistics of the start-point and end-point distances and angle differences between the teach path and the repeat paths. The average and standard deviation in meters and degrees are reported for both forward and backward repeats

VTR Variation of	avg. $\pm$ std. dev.
Start-Point Distances-Forward Motion	$2.65 \pm 1.34$ (m)
End-Point Distances-Forward Motion	$0.41 \pm 0.26$ (m)
Start-Point Distances-Backward Motion	$1.94 \pm 1.40$ (m)
End-Point Distances-Backward Motion	$0.60 \pm 0.20$ (m)
Starting Angle Differences-Forward Motion	$155.87^\circ \pm 125.94^\circ$
Ending Angle Differences-Forward Motion	$10.50^\circ \pm 9.40^\circ$
Starting Angle Differences-Backward Motion	$149.85^\circ \pm 76.69^\circ$
Ending Angle Differences-Backward Motion	$63.09^\circ \pm 19.15^\circ$

and standard deviation in meters are given in the Table I. The end-point statistics were also computed. As we can see, for forward repeats, the end points were on average almost one robot diameter (0.41 m) away from the teach path end point, despite the large range and variance of the start points in repeats. Also, the average and standard deviation of starting and ending robot angles for forward and backward repeats compared to the teach phase are given in the Table I. The algorithm has been able to move the robot to the ending point with an average of 10 degrees of error with respect to the teach phase. It is important to mention, due to the scale ambiguity of generated map (between each generated semantic map and the real world) we cannot report object localization errors. Also, our focus was mainly on pose independence, that is, being able to repeat the teach path on a relatively accurate path and end up at a point closed to the teach path end-point, starting from a variety of poses. Therefore, we do not attempt to track the teach path as close as possible, which can increase the motion smoothness of the repeat path produced by our algorithm.

We compare the performance of our algorithm to two recent closely-related VTR algorithms, [19] and [25], using the same teach path in the same testing environment. The robot was placed in various locations with similar initial poses as some of the tested poses for our own algorithm shown in figure 6. Figure 7 shows the performance of these two algorithms. In most cases, they cannot follow the teach path as well as our algorithm as shown in figure 6. The VTR algorithm from [19] also uses semantic objects to try to find the correct repeat path, similar to our work. However, our tests show that the robot often loses track of the teach path before it reaches the end point. Sometimes, the robot can even get stuck near the initial position and is not able to finish the test. The VTR algorithm from [25] mainly relies on recorded odometry data and tries to reduce the robot pose error using the captured image features. Not surprisingly, the robot cannot find similar image features to repeat the teach path, when there is a drastic change in the viewing angle. We do not report starting and ending position and angle errors for these methods in Table I, since they are unable to finish many tests.

### B. Backward Repeat

We evaluate the VTR algorithm in the backward direction with the same experimental setup as in the forward repeats. The only difference is that the robot would move toward the starting point instead of the end point on the teach path. Figure 8 shows the seven backward repeat motions tested in

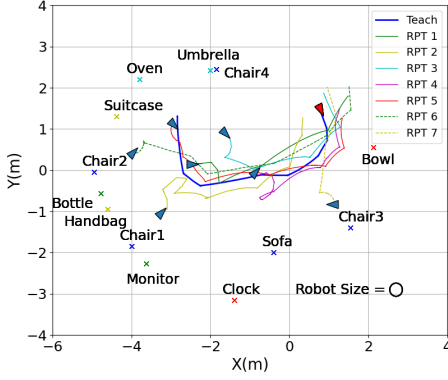


Fig. 8. The teach and seven repeat motions evaluated by Vicon motion capture system for backward repeat tests. Triangle arrows indicate the initial poses of the robot, which can vary significantly such as a closed to  $180^\circ$  viewing angle difference.

the lab. Again, we report the average and standard deviation of start-point and end-point distances and angles with respect to the teach path in Table I. The results indicate reasonable accuracy of our algorithm in backward repeats. In most cases, the path was followed in reasonable accuracy with respect to the size of the robot, especially considering the significant viewing angle differences, sometimes close to  $180^\circ$ , between the teach path and the backward repeats.

### C. Robustness Towards Environmental Changes

We also evaluate the robustness of the proposed algorithm to occasional environmental changes. As a test case, we randomly moved five of the objects, i.e., a chair, handbag, clock, oven, and the suitcase that were semantic landmarks in the teach map, to the new places. Seven more forward repeat tests were performed in the changed setting. Figure 9 shows the teach and repeat paths as well as the new positions of all landmarks. The repeat motions were as accurate as before qualitatively, as can be seen in figure 9, and quantitatively too, achieving comparable performance as reported in Table I. Similarly, other objects could be moved or removed from the environment without degrading the repeat path. Our VTR is robust to multiple environmental changes as long as the number of

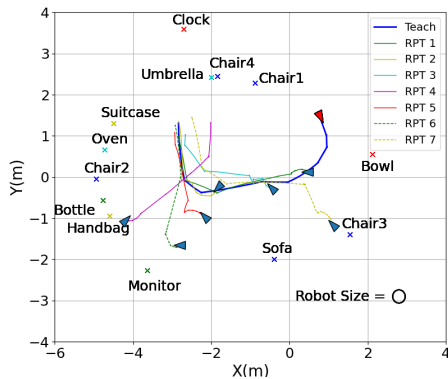


Fig. 9. The teach and repeat motions captured by our Vicon motion capture system, after relocation of five objects. Old location of each object can be seen in figure 6 and figure 8.

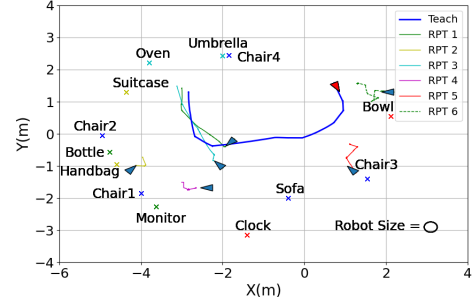


Fig. 10. The teach and repeat motions captured by our Vicon motion capture system, while purely using ORB-SLAM for camera pose relocalization.

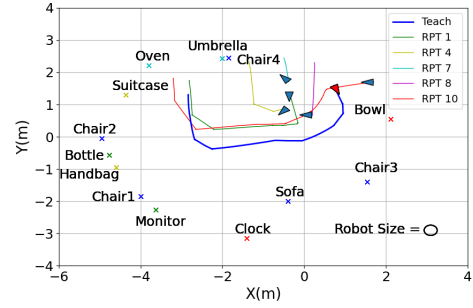


Fig. 11. The teach and repeat motions captured by our Vicon motion capture system, after relocalizing the robot between teach and repeat initial points by combining Superglue and RANSAC algorithms.

relocated objects is less than half of all the semantic objects, as our algorithm only uses the top half of objects with the lowest positional error to relocalize as explained in section V-A. This algorithm feature holds for backward repeats as well.

### D. Comparison with Image-Based Algorithms

In order to demonstrate the advantages of our proposed algorithm, we compared our method with two image-based localization algorithms. In the first case, we tested the relocalization capability of only using ORB-SLAM [6] as a visual SLAM method. Previously, reference [5] had shown ORB-SLAM is not capable of relocalizing the camera inside a created map beyond 30 degrees due to the lack of feature matches. To validate this again, we moved the robot through the teach path while using ORB-SLAM to create a map. Then, we moved the robot to each one of the repeat paths initial points and slightly moved and rotated the robot to re-initialize the ORB-SLAM. In most of the cases, where the camera pose was highly different with respect to the teach path, the ORB-SLAM was not able to relocalize the camera. However, in some of the cases with similar camera poses between teach and repeat phases, the ORB-SLAM could relocalize the camera position and the robot moved toward the end-point, figure 10.

Also, we compared our method with a relocalization algorithm combining Superglue [26] and RANSAC [27]. In order to help the algorithm, we rotated the camera toward the area that the camera is looking at the teach phase initial pose. Therefore, there would be more similarity between teach and repeat phases initial poses. After calculating the transformation



by superglue+RANSAC, we moved the robot to the estimated point and performed the repeat motion. The robot used the odometry data to move toward the end-point. We noticed this relocalization algorithm works accurately when the teach and repeat phases initial points are not far from each other. Figure 11 shows some of the repeat motions after relocalizing by superglue+RANSAC. As one can see, the repeat initial points that are close to the teach initial point, i.e., Repeat1 and Repeat10, could repeat the teach path accurately. But the other points performed the repeat motion poorly. The equivalent point for each repeat initial point and path can be found in figure 6. We omit these methods from Table I since they do not complete many of the tests.

### E. Memory Requirement and Computation Time Comparisons

We measured the memory requirement of our algorithm and compared it with [19], [25] and pure ORB-SLAM in similar conditions. We measured the peak memory usage and computation time of the whole algorithm. Our method never used more than 1.6 GB of memory. This number was 1.8 GB for [19] and 40 GB for [25] as they store all the images of the teach path. Also, the Superglue+RANSAC method required 1.7 GB. Pure ORB-SLAM required 0.6 GB for the cost of much lower performance as figure 10 shows. Compared to ORB-SLAM, our method requires more memory for the object detection module, which is required for the robustness and pose-independence of the VTR algorithm. We also measured the maximum computation time, over all execution time steps, of each method mmilitwith similar image resolution and frame rates. References [19], [25], Superglue+RANSAC, pure ORB-SLAM and our method required 0.001, 0.1, 0.1, 0.003 and 0.03 seconds respectively. Our python implementation is not optimized, but still fast enough for running our algorithm in real-time.

## VII. CONCLUSION, DISCUSSION AND FUTURE WORK

We have proposed a novel VTR algorithm using 3D semantic maps. The algorithm is robust to large robot starting pose changes in the repeat phase. Our key insight is to build 3D semantic maps, containing semantic labels and positions of the objects in the environment as well as camera poses, during the teach phase. In the repeat phase, the algorithm relocalizes the robot in the new map based on the found objects. We tested our method in both forward and backward modes starting from various locations inside the lab. Our algorithm demonstrated robustness toward significant starting pose variations, as well as tolerance to environmental changes. Our algorithm compares favorably with respect to two other most-relevant recent VTR methods and two image-based relocalization algorithms in terms of both accuracy and robustness.

Although we were able to show a simple semantic map can help to relocalize the robot and perform an accurate, robust, and pose-independent VTR, there are still multiple areas for future research. For example, we showed the performance of our algorithm in an indoor area. While it was not the main focus of this paper, it would be interesting to test our algorithm for repeating for much longer and possibly outdoor paths. Also, the accuracy of the semantic maps could be improved by a more sophisticated 3D position estimation algorithm. In certain environments with very few objects, it may be useful to utilize

repetitive objects in the environment in the future. We also wish to improve the end-point accuracy for applications that require return trips and loops, and consider environments with other mobile agents.

## REFERENCES

- [1] D. G. Lowe, "Object recognition from local scale-invariant features," *ICCV*, vol. 2, 1999.
- [2] H. Bay, T. Tuytelaars and L. Van Gool, "Surf: Speeded up robust features," *ECCV*, Springer, 2006.
- [3] M. Calonder, V. Lepetit, C. Strecha and P. Fua, "Brief: Binary robust independent elementary features," *ECCV*, Springer, 2010.
- [4] E. Rublee, V. Rabaud, K. Konolige and G. Bradski, "ORB: An efficient alternative to SIFT or SURF," *ICCV*, IEEE, 2011.
- [5] J. Lee, D. Meger and G. Dudek, "Semantic mapping for view-invariant relocalization," *ICRA*, IEEE, 2019.
- [6] R. Mur-Artal and J. D. Tardós, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *Transactions on robotics*, IEEE, 2017, pp. 1255-1262.
- [7] G. Yu and J. M. Morel, "ASIFT: An algorithm for fully affine invariant comparison," *Image Processing On Line*, vol. 1, 2011, pp. 11-38.
- [8] A. Pfunder, A. P. Schoelling and T. D. Barfoot, "A proof-of-concept demonstration of visual teach and repeat on a quadcopter using an altitude sensor and a monocular camera," *CRV*, IEEE, 2014.
- [9] T. Nguyen, G. K. Mann, R. G. Gosine and A. Vardy, "Appearance-based visual-teach-and-repeat navigation technique for micro aerial vehicle," *Journal of Intelligent & Robotic Systems*, Springer, vol. 84, no. 1, pp. 217-240, 2016.
- [10] S. Milz, G. Arbeiter, C. Witt, B. Abdallah, and S. Yogamani, "Visual slam for automated driving: Exploring the applications of deep learning," *CVPRW*, IEEE, 2018.
- [11] P. Furgale and T. D. Barfoot, "Visual teach and repeat for long-range rover autonomy," *Journal of Field Robotics*, vol. 27, no. 5, 2010.
- [12] M. Paton, K. MacTavish, M. Warren and T. D. Barfoot, "Bridging the appearance gap: Multi-experience localization for long-term visual teach and repeat," *IROS*, IEEE, 2016.
- [13] C. J. Ostafew, A. P. Schoelling and T. D. Barfoot, "Visual teach and repeat, repeat, repeat: Iterative learning control to improve mobile robot path tracking in challenging outdoor environments," *IROS*, IEEE, 2013.
- [14] L. Celement, J. Kelly and T. D. Barfoot, "Robust monocular visual teach and repeat aided by local ground planarity and color-constant imagery," *Journal of field robotics*, vol. 34, no. 1, pp. 74-97, 2017.
- [15] M. Warren, M. Greeff, B. Patel and J. Collier, A. Schoelling, T. D. Barfoot, "There's no place like home: Visual teach and repeat for emergency return of multirotor uavs during gps failure," *RA-Letter*, vol. 4, no. 1, pp. 161-168, 2018.
- [16] M. Gridseth and T. D. Barfoot, "Towards direct localization for visual teach and repeat," *CRV*, IEEE, 2019.
- [17] M. Gridseth and T. D. Barfoot, "DeepMEL: Compiling Visual Multi-Experience Localization into a Deep Neural Network," *ICRA*, IEEE, 2020.
- [18] L. G. Camara, T. Pivoňka, M. Jílek, C. Gäbert, K. Košnar and L. Přeucil, "Accurate and robust teach and repeat navigation by visual place recognition: A cnn approach," *IROS*, IEEE, 2020.
- [19] A. G. Toudeshki, F. Shamshirdar and R. Vaughan, "Robust UAV visual teach and repeat using only sparse semantic object features," *CRV*, IEEE, 2018.
- [20] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv:1804.02767*, 2018.
- [21] M. Bjelonic, "YOLO ROS: Real-Time Object Detection for ROS," <https://github.com/leggedrobotics/darknet-ros>, 2018.
- [22] D. Singh, E. Trivedi, Y. Sharma and V. Niranjan, "TurtleBot: Design and Hardware Component Selection," *GUCon*, IEEE, 2018.
- [23] S. Ren, K. He, R. Girshick and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *Advances in neural information processing systems*, vol. 28, pp. 91-99, 2015.
- [24] S. Y. Bao, M. Bagra, Y. Chao and S. Savarese, "Semantic structure from motion with points, regions, and objects," *CVPR*, IEEE, 2012.
- [25] D. Dall'Osto, T. Fischer and M. Milford, "Fast and robust bio-inspired teach and repeat navigation," *IROS*, IEEE, 2021.
- [26] P. Sarling, D. DeTone, T. Malisiewicz and A. Rabinovich, "Superglue: Learning feature matching with graph neural networks," *CVPR*, IEEE, 2020.
- [27] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, IEEE, vol. 24, no. 6, pp. 381-395, 1981.
- [28] F. Shamshirdar, "Deep neural networks for robot navigation using visual teach and repeat," PhD diss., *Applied Sciences: School of Computing Science*, Simon Fraser University, 2019.