

Learning to Prevent Monocular SLAM Failure using Reinforcement Learning

Vignesh Prasad
Embedded Systems & Robotics
TCS Research & Innovation
Kolkata, India
vignesh.prasad@tcs.com

Karmesh Yadav
Robotics Institute
Carnegie Mellon University
Pittsburgh, USA
ykarmesh@gmail.com

Rohitashva Singh Saurabh
Dept. of Robotics Engineering
John Hopkins University
Baltimore, USA
saurabhiitg11@gmail.com

Swapnil Daga
Robotics Research Center
KCIS, IIIT Hyderabad, India
swapnil.daga@research.iiit.ac.in

Nahas Pareekutty
Robotics Research Center
KCIS, IIIT Hyderabad, India
nahas.p@research.iiit.ac.in

K. Madhava Krishna
Robotics Research Center
KCIS, IIIT Hyderabad, India
mkrishna@iiit.ac.in

Balaraman Ravindran
Dept. of Computer Science,
Indian Institute of Technology
Madras, India
ravi@cse.iitm.ac.in

Brojeshwar Bhowmick
Embedded Systems & Robotics
TCS Research & Innovation
Kolkata, India
b.bhowmick@tcs.com

ABSTRACT

Monocular SLAM refers to using a single camera to estimate robot ego motion while building a map of the environment. While Monocular SLAM is a well studied problem, automating Monocular SLAM by integrating it with trajectory planning frameworks is particularly challenging. This paper presents a novel formulation based on Reinforcement Learning (RL) that generates fail safe trajectories wherein the SLAM generated outputs do not deviate largely from their true values. Quintessentially, the RL framework successfully learns the otherwise complex relation between perceptual inputs and motor actions and uses this knowledge to generate trajectories that do not cause failure of SLAM. We show systematically in simulations how the quality of the SLAM dramatically improves when trajectories are computed using RL. Our method scales effectively across Monocular SLAM frameworks in both simulation and in real world experiments with a mobile robot.

CCS CONCEPTS

• **Computing methodologies** → **Motion path planning; Reinforcement learning**; • **Computer systems organization** → *Robotics; Robotic autonomy*.

KEYWORDS

Reinforcement Learning, Robot Navigation, Monocular SLAM

ACM Reference Format:

Vignesh Prasad, Karmesh Yadav, Rohitashva Singh Saurabh, Swapnil Daga, Nahas Pareekutty, K. Madhava Krishna, Balaraman Ravindran, and Brojeshwar Bhowmick. 2018. Learning to Prevent Monocular SLAM Failure using Reinforcement Learning. In *11th Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP 2018), December 18–22, 2018, Hyderabad, India*, Anoop M. Nambodiri, Vineeth Balasubramanian, Amit Roy-Chowdhury, and Guido Gerig (Eds.). ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3293353.3293400>

1 INTRODUCTION

Simultaneous Planning Localization and Mapping (SPLAM) or Active SLAM has been a popular area of research over the years. The main theme is to compute a set of control actions for the mobile robot such that either the uncertainty of the robot or the combined uncertainty of robot and map states are bounded. Traditionally, SPLAM frameworks have been approached through either Model Predictive Control [14, 15] or Information Gain [12] paradigms. In recent times, Belief Space planning paradigms have tended to compute the control law in continuous domain [2, 8] taking a leaf out of SLAM frameworks that model it as a least squares problem [9]. Planning under uncertainty using monocular vision by considering scene texture in addition to the geometric structure of the scene has also shown to improve the localization uncertainty [4].

However, all the above works assume the availability of immediate range data, map uncertainty estimates or dense depth maps of the scene. In contrast, Monocular SLAM gives sparse depth estimates, which are highly inaccurate in texture-less or low texture scenes. These cases usually cause a failure in the feature tracking and consequently in SLAM itself. Moreover, planar scenes and in-place rotations lead to degeneracies in camera pose estimates. These combined with the non-linear nature of the camera projection operation lead to complexities in Monocular SLAM estimates.

Authors Vignesh and Karmesh contributed equally.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICVGIP 2018, December 18–22, 2018, Hyderabad, India

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-6615-1/18/12... \$15.00

<https://doi.org/10.1145/3293353.3293400>

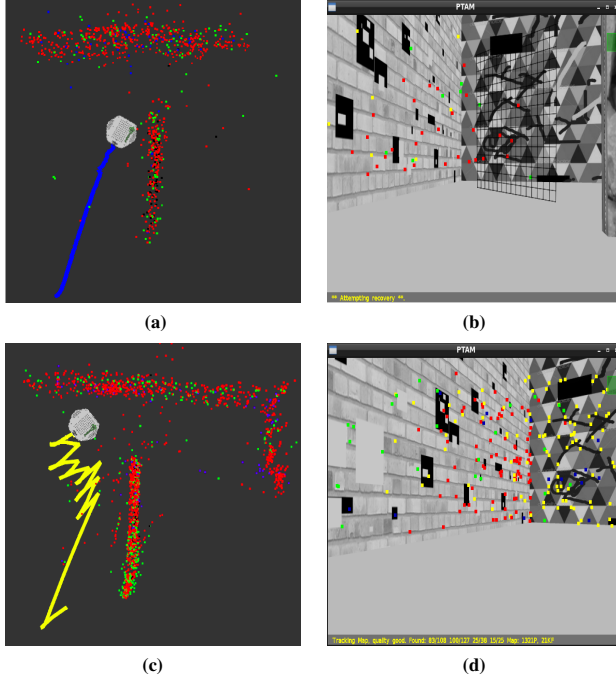


Figure 1: Robot Trajectory estimate and Map (a) with a naive uninformed planner and (c) with the RL-based planner. The corresponding Monocular SLAM(PTAM) feedback is shown in (b) and (d) respectively.

Literature seems sparse on SPLAM in a monocular setting. A Next-Best-View (NBV) approach [16], based on estimated measures of the localization quality, showed good results for goal based trajectory planning as well as strategic exploration on Micro Aerial Vehicles (MAVs). Nonetheless it is shown subsequently that the current method evaluates better on various maps in terms of SLAM failures. Autonomous navigation with a downward facing camera for Micro Helicopters was shown in [25]. In our case, we try to autonomously navigate a non-holonomic robots with a front facing camera. This is more constrained than the above mentioned works where the use of MAVs allow smooth lateral motions thus providing wide baselines to improve the quality of Monocular SLAM, whereas in our case we use a ground robot restricted to move either forward or backward.

Prasad *et al.* [19] approach the problem using an Apprenticeship Learning formulation [1]. However their approach is time consuming during training, where they evaluate the learnt reward at each iteration and the efficacy of this formulation for larger scale maps is also unproven.

In contrast, this paper formulates the SPLAM problem through a Reinforcement Learning paradigm in a map agnostic fashion. Our approach is trained on a single map and is applied on a variety of maps without any re-training/fine-tuning. Instead of learning outright an optimal policy, the current framework learns actions detrimental to SLAM using Temporal Difference Learning methods [22]. We learn an optimal action-value function with respect to SLAM failure which is used to filter out potentially unsafe actions. This learnt optimal Action Filter is seamlessly used across a variety of maps showing significant reduction in Monocular SLAM failures, which

are mainly characterized by tracking failures or due to large errors in the pose priors used for estimation.

By using RL in an episodic manner, the result of each failure is propagated back to earlier states over multiple episodes, making the predictions of future failures more likely in the earlier stages once the learning converges. Analysis shows significant reduction in SLAM failures when the learned Q values are used to filter actions in trajectories generated by routine methods such as sampling based planners [13] or trajectory optimization routines [7] [21]. Qualitative and quantitative analysis and comparisons with supervised learning approaches, 3D point overlap maximization techniques, and predicting Localization Quality estimates, described in [16], showcase the superior performance of the proposed method. The effectiveness of RL based navigation is also demonstrated with different feature based Monocular SLAM systems, in both simulated and real world indoor scenes with a ground robot. We further show the advantage of our algorithm when augmented with relocalization after SLAM failure and with multi-modal sensor fusion with SLAM, leading to almost 100% prevention of SLAM failure.

The keynote is that almost all previous Monocular SLAM results have been shown with a handheld camera or on a robot that is teleoperated by a user. In such scenarios, human intuition and experience is often responsible for the success of SLAM. For example, in the case of handheld cameras it is common for the user to repeatedly scan the same area to improve the quality of results. In this paper the robot negotiates turns through a sequence of back and forth motions, tacitly avoiding a single shot acute turn or bend. RL based methods are amenable to learn such experience and intuition and we dovetail it to a RL based planner forming the essential theme of this work. This is precisely captured in the trajectory of Fig. 1c, where a sharp turn is decomposed into a sequence of back and forth maneuvers that has been learned by RL. Such a maneuver provides for SLAM stability by attenuating a sharp turn and allowing the vehicle to look at an area for a longer time. This helps to ensure persistence of mapped points over multiple frames thereby improving the accuracy of the estimates of both the pose and the sparse map.

2 BACKGROUND

2.1 Monocular SLAM

Monocular Simultaneous Localization and Mapping (Monocular SLAM) refers to the method of estimating the camera pose while performing a 3D reconstruction of the environment using only a monocular camera. **PTAM (Parallel Tracking and Mapping)** [10] was one of the first such SLAM systems. **ORB-SLAM** [17] is a newer feature based Monocular SLAM that builds on the main ideas of PTAM, using different image features for tracking as compared to PTAM and incorporates real time loop closure as well. **LSD-SLAM** [5] is another state-of-the-art Monocular SLAM system which works with image intensities and creates a semi-dense map of the environment. A survey of various Monocular SLAM methods can be found in [26].

While these methods have shown effective SLAM capabilities, they are susceptible to erroneous pose estimates due to insufficient tracking of mapped points or motion induced errors such as large or in-place rotations. The latter can be mitigated to an extent by restricting robot motion, but this may lead to inefficiencies in navigation.

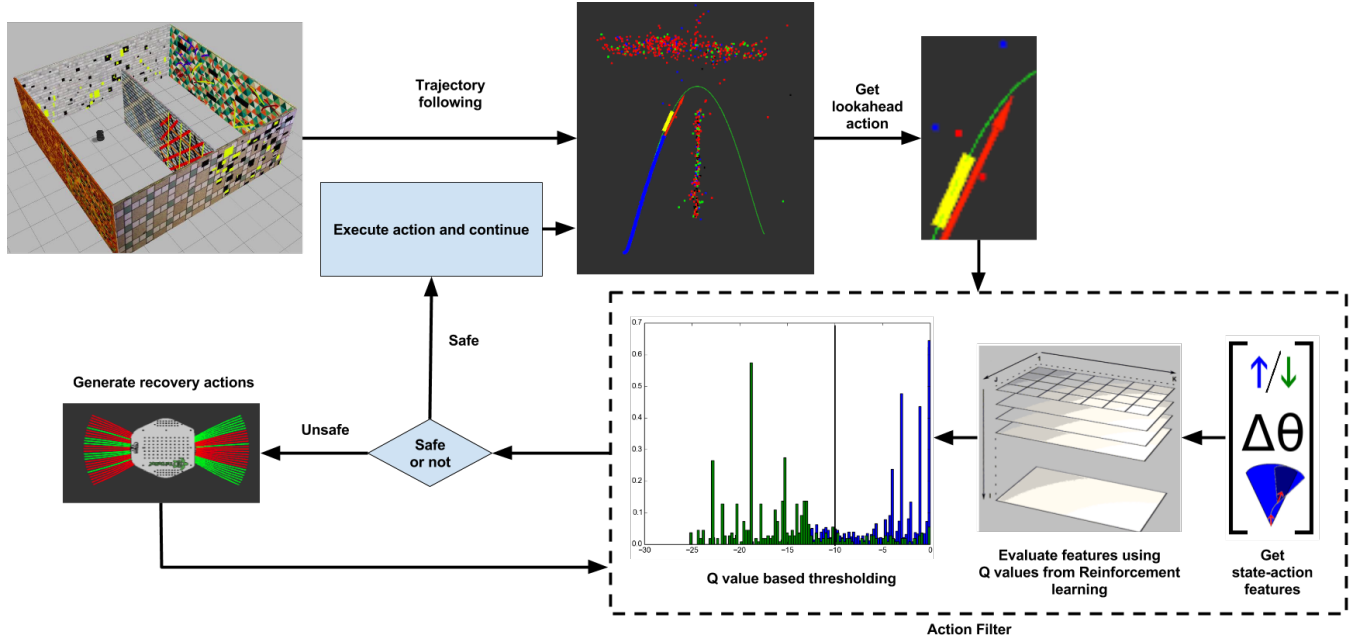


Figure 2: Overview of the RL based Navigational Planner. Once a global trajectory is generated by an underlying planner, as it is executed, the subsequent part of the trajectory (shown with the yellow marker), along with the current set of points are taken as the state-action pair and the Q value for it is calculated. The Q value is checked to see if the motion that is to be executed is safe or not. If its not, then sample recovery actions are generated and among those with a safe Q value (shown in green in the bottom left), the one that best aligns the robot with the global trajectory is executed. After this, a new global trajectory is generated from the new pose and this cycle continues till the robot reaches the goal.

A promising solution would therefore be based on automatically learning robot behavior that navigates while keeping SLAM failures to a minimum. Considering the requirements of self-learning and adaptability, Reinforcement Learning forms a good candidate for such a solution.

2.2 Reinforcement Learning

Reinforcement Learning (RL) is a learning method based on Markov Decision Processes (MDPs) where actions are performed based on the current state of the system and rewards are obtained accordingly. The correspondence between states and the actions performed from each state is determined by the policy, π . The aim of the decision process is to optimize the policy, which in turn is performed by maximizing the return from each state. In formal notations, an MDP is a tuple (S, A, T, γ, R) , where

- S is a set of states.
- A is a set of actions.
- T is a set of state transition probabilities.
- γ is the discount factor. $\gamma \in [0, 1]$
- R is the reward function. $R : S \times A \rightarrow \mathbb{R}$

The effectiveness of an action is defined by the Q value of the action and the state from which it is performed i.e. state-action pair, denoted by $Q(s, a)$. Once learned effectively, it takes into account not just the immediate reward but also the future discounted rewards of the transitions that follow.

In this paper, we use Q-Learning [23, 24] for learning navigational strategies. Q-Learning is a widely used Temporal Difference

prediction method for control problems. It learns optimal Q values for a given MDP. Starting with an initial estimate, these values are updated incrementally through learning episodes or samples, in a way similar to dynamic programming. The update for performing an action a from a state s , reaching a new state s' and obtaining a reward r is as follows:

$$Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a') - Q(s, a)] \quad (1)$$

where a' is the action that would be performed from the next state s' , α is the learning rate and γ is the discount factor. Once the Q values are learnt, the optimal action is the one which maximizes the Q value from the current state.

3 PREVENTING SLAM FAILURE USING RL

The motivation for this paper is to devise a navigational methodology for a robot that prevents Monocular SLAM failure. This is ensured by using RL to learn the relationship between robot actions and SLAM failure. The robot essentially learns to avoid potentially unsafe movements that can lead to SLAM failure. Fig. 2 shows a flowchart of the complete navigational planner, which is explained below in detail.

3.1 Navigation Overview

Initially, a trajectory is planned from the current location to a goal, as depicted by the green curve in the middle image of Fig. 2. Goals are selected at frontier locations of the map with enough overlap with already mapped points to prevent SLAM failure.

Once the robot starts moving, the immediate subsequent part of the trajectory, which is shown in yellow in the same image of Fig. 2, is parametrized into a state-action pair, details of which can be found in Sec. 3.2. An Action Filter evaluates whether this action will cause a SLAM failure or not. This is done by comparing the Q value of the corresponding state-action pair against a safety threshold. This is shown in the bottom right image of Fig. 2. If the Q value falls below the threshold, i.e. if an unsafe motion is predicted, the robot is stopped and potential recovery goals are generated and checked by the Action Filter. The recovery actions are trajectories to points at a distance of 1m from the current position in both forward and backward directions with $0^\circ - 30^\circ$ deviations on either side (left/right). This can be seen in the bottom left image of Fig. 2, where the green trajectories are the safe trajectories and the red ones are unsafe. We restrict the angle to 30° as angles larger than this, in such a small distance, would be detrimental for SLAM. We choose these actions at a distance of 1m as points further away would lead to lesser overlap between images and more importantly, in 3D scene points.

Among those predicted as safe, the pose that best aligns it with the global path is chosen for execution. If the chosen pose is too close to the unsafe ones, it is discarded and the pose with the best Q value is chosen. Once this recovery trajectory is executed, a new path is generated from the new location to the goal. This process is cyclically repeated until the goal is reached.

3.2 Learning to prevent SLAM failure using RL

To effectively design the Action Filter, the quality of an action should be primarily dependent on the probability of it causing a SLAM failure while being independent of the map(s) being used for learning.

The objective of the RL method is to evaluate the quality of an action performed from a state with respect to the SLAM pose estimate. The model was trained in a simulated environment, by generating sample paths from the robot's current pose. We use Q-Learning for performing the learning, which takes as input a state-action pair (s, a) . This leads us to use the SLAM related outputs from the current and the next pose. Our state space comprises of the number of common map points seen between the subsequent poses, N_{3D} while the action space includes the direction of motion (forward or backward), represented as η and the angle change, $\Delta\theta$ between the current and the next pose.

We observed an increasing trend in the frequency of failures as the change in heading angle increased. This causes the pose estimate of SLAM to deteriorate, which leads to SLAM failure. Therefore we restrict our learning to a maximum of 30° . Coming to N_{3D} , we keep a cap of 600 points as anything more would anyway not contribute to SLAM failure or cause deterioration of SLAM estimates.

We use a lookup table representation for the Q function, by discretizing the parametrized state-action pair. Details of this discretization are given in table 1.

The Q values were learnt by performing Q-Learning with a weighted combination of the above parameters as the reward function. The reward function weights are chosen by making a few assumptions about the way Monocular SLAM Systems behave. Other than observing the change in heading angle and overlap of points,

Table 1: Discretization of State-Action pairs

Parameter	Range	No. of discretized values
η	forward or backward	2
$\Delta\theta$	0° to 30°	20
N_{3D}	0 - 600	20

the status of SLAM after performing an action is also used, which we denote as Φ . This is a Boolean value which tells us if SLAM failure occurs due to a particular action. This is an observable entity obtained from Monocular SLAM used during the training phase only.

The reward function of our model is a weighted combination of these parameters given in Eq. 2.

$$R(s, a) = \omega_{3D}N_{3D} + \omega_\theta\Delta\theta + \omega_\Phi\Phi + \omega_0 \quad (2)$$

The weights were assigned based on the extent to which these parameters affect Monocular SLAM.

- The weight ω_{3D} is given a positive value of 0.0167 (with the maximum value of N_{3D} being 600, it allows a maximum reward of +10) as the larger the overlap between frames, the better it is for Monocular SLAM.
- For ω_θ , a negative weight of -0.1 is given (with the lowest reward being -3) as larger angle change degrades the quality of SLAM.
- Since our main aim is to prevent SLAM failure a relatively large value of -10 is given to ω_Φ .
- The value of the bias term, ω_0 is set to -10 so that the overall reward is kept negative.
- There is no separate reward for successfully reaching a goal state. Our aim is not to optimally reach the goal location, but to prevent SLAM failure along the way.

4 EXPERIMENTATION

4.1 Training

Gazebo [11] is a framework for robot simulations and it reproduces dynamic environments like those a robot would encounter in the real world. Indoor maps were created in Gazebo to simulate real-life environments.

Training is carried out in a training map (7mx10m) using PTAM, shown in Fig. 3. Navigation during episodic learning consisted of random walks with an Action Filter derived from the learnt Q values. An episode is the sequence of steps until SLAM failure. Initially, motions were randomly generated to learn an initial set of Q values. Once the Q values are initialized, we use an incremental ϵ -Greedy policy to reinforce the learnt Q values after each episode using Q-Learning. After each set of episodes, the value of ϵ is incremented by a small amount until it reaches 0.9, as can be seen in Alg. 1. Basically, once a basic set of values are learnt, we keep increasing the exploitation ratio in order to further reinforce what is learnt. This can be seen in the increasing trend in the number of steps taken till SLAM failure, as seen in Fig. 4. Initially, due to the high level of randomness, there is a high variance in the output but as the learning progresses, there is an increasing trend in the average number of steps taken till SLAM failure.




Figure 3: A map of the simulated training environment in Gazebo.

One thing to note is that unlike most RL problems where the average reward per step or similar metrics would plateau, our metric, i.e. the average number of steps to SLAM failure, would keep increasing as the learning progresses. Ideally, the robot would never get into a situation where SLAM would fail. One possible way to combat this is to keep a cap on the maximum number of steps per episode. This could possibly help speed the learning as the larger the number of steps go, the greater they get discounted so it would take longer time to propagate its effect to the initial states. However, we defer this exercise for our future work.




Figure 4: The average steps to taken by the robot until SLAM failure in the training phase.

Once the Q values are learnt, we use it in the Action Filter to classify actions as failing or not. The optimal quality threshold for the Action Filter would precisely classify state-action pairs into SLAM failures/non-failures. The quality threshold was chosen based on the primary objective of minimizing false positives (unsafe actions classified as safe) and a secondary objective of minimizing false negatives (safe actions classified as unsafe). From our experiments, we found a threshold of -10 to be suitable for achieving good performance.

The main takeaway is that the training was performed on a single map and was seamlessly applied on various maps of varying structure

Algorithm 1 Incremental ϵ -Greedy Training

```

 $\epsilon \leftarrow 0.1$ 
Repeat (while  $\epsilon < 1$ ):
  Repeat (for  $N$  episodes):
    Initialize state  $s$ 
    Repeat (for each step in the episode):
      Choose action  $a$  at  $s$  using an  $\epsilon$ -Greedy policy
      Execute action  $a$  and get reward  $r$  and reach next state  $s'$ 
       $Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$ 
       $s \leftarrow s'$ 
     $\epsilon \leftarrow \epsilon + 0.1$ 

```

and sizes, without the need for retraining. The same Q-values learnt using PTAM were able to be applied directly to ORB-SLAM as well, by just changing the thresholds. This highlights the effectiveness of our map-agnostic approach and also shows that it can be effortlessly applied to different Monocular SLAM algorithms as well.

4.2 Comparison with Alternative Policies

In order to get further insight into the effectiveness of our method, we substituted the learnt policy with three alternative policies. One using supervised learning, one based on the common mapped points between subsequent poses, which we defined earlier as N_{3D} and a Next-Best-View method.

We decided to use support vector machines (SVM) [3] over other classifiers like artificial neural networks, decision trees etc. due to better average performance tested using Monte Carlo cross-validation. A SVM with a radial basis function (rbf) kernel, implemented using scikit-learn [18], is trained to classify trajectories as failing or not. The distance of samples from the decision boundary are used to give a score to the trajectories. The same parameters which were used by RL, are used as inputs for the SVM.

Superior performance is also shown in comparison to directly using N_{3D} to detect failures and choose actions. This proves that feature overlap isn't the only parameter adversely affecting Monocular SLAM and other parameters need to be considered to effectively navigate in such a setting.

Comparison with a Next-Best-View formulation [16], based on estimating the localization quality is also shown. This is done based on the Geometric Point Quality and the Point Recognition Probability of the mapped points. The Geometric Point Quality penalizes points with low triangulation angles, as this increases localization uncertainty. The Point Recognition Probability gives an estimate of the recognizability of a point based on the viewing angle and the scale at which the point was originally found. This provides a faster evaluation of map uncertainty as compared to calculating covariance matrices. A detailed explanation can be found in [16].

5 RESULT OF RL BASED NAVIGATION

Experiments were conducted by navigating a Turtlebot on different maps in Gazebo environment. For the physical implementation a Pioneer 3-DX was used. In all the experiments, the color image output of a Microsoft Kinect was used as the video source. Experiments were carried out on a laptop with Intel Core i7-7700U 2.80GHz CPU

running Ubuntu 16.04 using Robot Operating System (ROS) [20] for controlling the robot and performing SLAM.

5.1 Steps till SLAM Failure




Figure 5: Comparison of average steps to SLAM failure using SVM, Localization Quality based Next-best-view (NBV)[16] and RL. The results are averaged over 10 episodes each.

To evaluate the performance of our method, we compare the average steps executed by the robot from SLAM initialization till SLAM failure. By step, we mean a trajectory to a distance of 1m from the current position in both forward and backward directions with $0^\circ - 30^\circ$ deviations on either side (left/right), similar to the recovery action generation. This is evaluated in the training map, shown in Fig. 3. The robot starts at an arbitrary location in the map and performs actions based on the policy to be evaluated.

As seen in Fig. 5, our policy performs better than all the other policies considered. Since the RL model is episodically trained, the result of each failure gets propagated back over episodes, making the predictions of future failures more likely in the earlier stages as the learning goes on. Hence the predictions are more accurate than an SVM, which was trained by just a single label for each step (failing or not).

5.2 Goal Based Navigation

In order to show the robustness of our approach over different trajectory generators, we combine our RL based approach with Bernstein curves [7] for trajectory generation and with the ROS Navigation Stack, in which navigation carried out in two levels where Informed RRT* [6] behaves as the global planner and Timed Elastic Band (TEB) [21] performs the role of the trajectory generator. We show results both without (naive) and with RL, to show that an informed approach using RL is more effective than an uninformed or naive path planner.

5.2.1 Bernstein based planning. After training the robot with PTAM, we test our approach in both simulation and real world. Navigation is performed using Bernstein Curves for path planning, with the robot moving from a predefined start to goal point on different maps using PTAM or ORB SLAM. The results obtained from the RL based planner are also compared against those obtained by using SVM, N_3D and NBV[16] based policies. Table 2 summarizes the

results of our experiments with each row showing the number of successful runs with the total number of runs given in the brackets. Maps 1-3 are of experiments done in Gazebo and map 4 is of a real environment run.

The last column of Table 2 shows the number of successful exploration run in a real world environment with Bernstein curves. The robot was made to enter a room from the adjoining corridor and explore the area in the room. The last row of table 2, corresponding to Map 4, shows the results for this run.

Table 2: Results for Goal Based Trajectory Planning (on Map 1-4) using Bernstein with PTAM and ORB SLAM in terms of no. of successful runs. The no. of runs on each map are given in brackets. These results can be seen in Fig. 6

SLAM	Policy	Map 1 (10 runs)	Map 2 (10 runs)	Map 3 (10 runs)	Map 4 (5 runs)
PTAM	Naive	2	0	0	1
PTAM	SVM	4	3	3	-
PTAM	N_3D	6	0	0	-
PTAM	NBV[16]	7	2	2	-
PTAM	RL	9	8	6	4
ORB-SLAM	Naive	3	0	1	-
ORB-SLAM	RL	9	8	6	5

5.2.2 Navigation Stack based planning. In Fig. 7 we show the simulation environments and full maps recovered with ORB SLAM using the RL based Planner based on the ROS Navigation Stack. Until now, autonomous navigation with monocular camera has been done only for short distances. We specifically designed these simulation environments to be difficult for navigation, with larger map sizes(20mX20m) and sparse feature scenes like featureless walls. Goals were given in different parts of the map and the robot had to explore its way to the goal without any prior knowledge of the surrounding. Table 3 shows the qualitative results obtained on Maps 5-7, where the RL based planner works considerably better than other methods. Exploration makes performing Monocular SLAM especially difficult due to the increased rotatory motion of the camera which could lead to distorted maps. Results of an exploratory run on a real robot where goals are user-specified are shown in Fig. 8. The robot is made to navigate from inside a room to the adjoining corridor. Such a transition in the environment is pretty challenging for simple Monocular SLAM.

Table 3: Number of Successful Runs (on Map 5-7) for Goal Based Navigation using Navigation Stack with ORB SLAM. These results can be seen in Fig. 7

Policy	Map 5 (15 runs)	Map 6 (15 runs)	Map 7 (15 runs)
Naive	3	2	2
RL	12	11	13
Naive + Odom	2	6	7
RL + Odom	14	13	14
Naive + Reloc + Odom	6	11	11
RL + Reloc + Odom	15	15	15

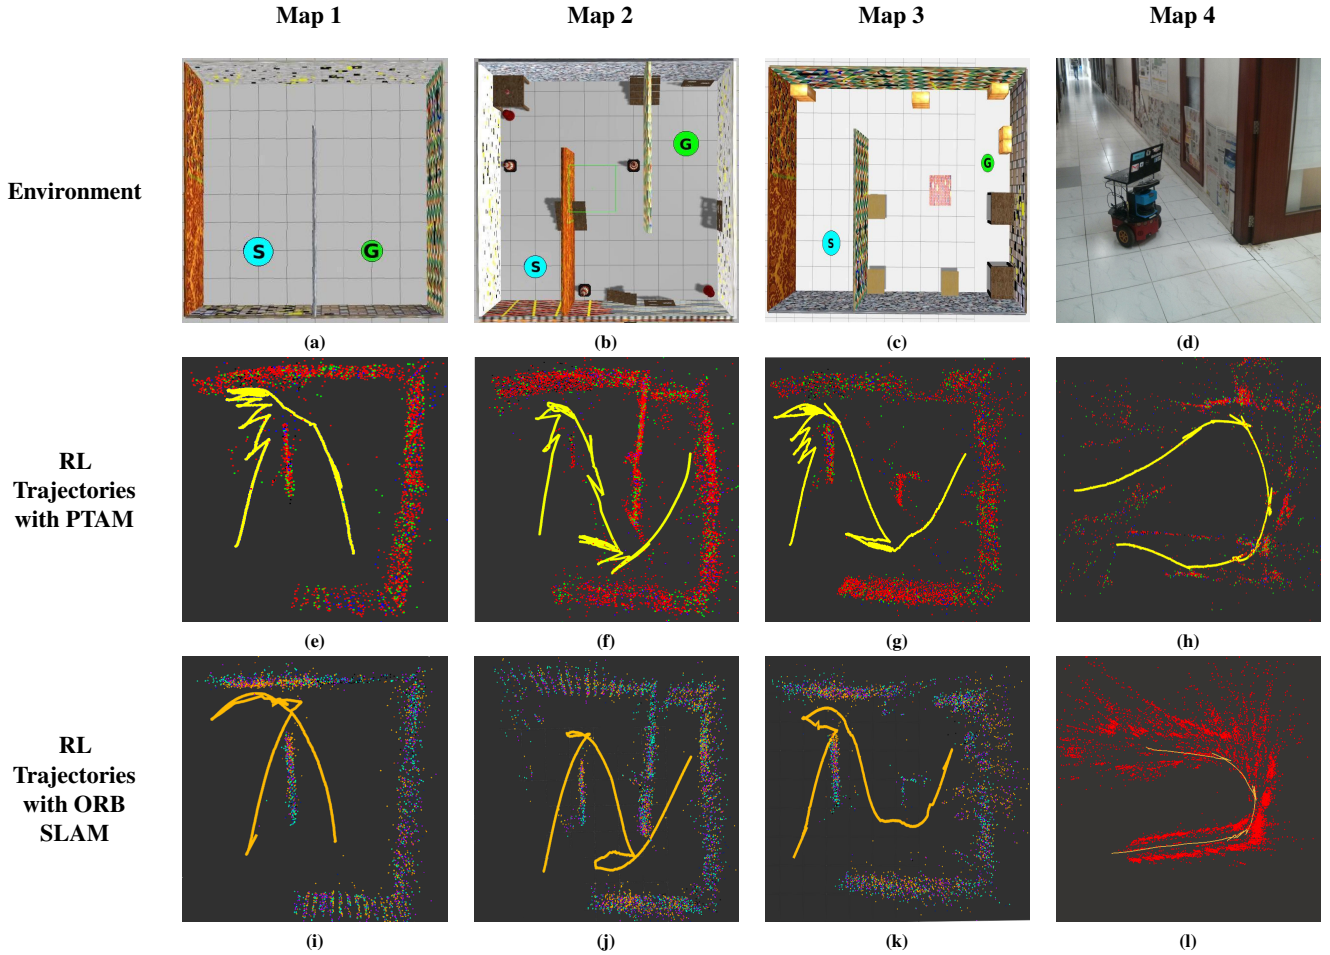


Figure 6: Results of Goal based trajectory navigation using Bernstein Curves. The Gazebo map shows the start and goal locations with blue and green circles, marked as "S" and "G" respectively. The second and third row shows the Trajectory and Map estimates using PTAM and ORB SLAM with the RL based planner respectively.

To demonstrate the advantage of our approach over methods which use sensor fusion for localization, we augment SLAM's localization output with wheel odometry data obtained from Gazebo. Multi-modal sensor fusion is expected to alleviate some of the localization issues during navigation. However, as visible from the 2nd and 3rd row of Table 3 (policies "RL" and "Naive + Odom" respectively), our RL based planner without wheel odometry performs better than a naive planner with sensor-fusion based localization output. Fusing our RL based Planner with the odometry is able to further improve the navigation capabilities as seen in the 4th row of Table 3 ("RL + Odom").

We also show the superiority of the RL based Planner against the ROS Navigation Stack which utilizes ORB SLAM's relocalization capabilities to restart SLAM after a tracking failure. When a failure occurs, the robot is manually given a goal facing a previously visible scene thus allowing relocalization. To allow the robot to keep moving after a tracking failure, the SLAM localization output is augmented with wheel odometry data. The 5th and 6th rows of Table 3 ("Naive + Reloc + Odom" and "RL + Reloc + Odom"), show the results of ORB SLAM with relocalization. We allow only 2 relocalization

attempts as increasing these attempts leads to degradation of the SLAM output while also increasing the time required to reach the goal. Finally, augmenting our RL based Planner with relocalization capabilities gives 100% success, with accurate reconstruction as indicated in the last row of Table 3.

6 CONCLUSION AND FUTURE WORK

We propose a map-agnostic Reinforcement Learning based navigational planner that learns to identify motions that can lead to Monocular SLAM failures and avoids them by performing recovery behaviors. The generated trajectories have minimal failures and pose estimation errors which is proved by carrying out experiments both in simulations and on real robots. This is different from traditional belief space planners or other SPLAM formulations due to their reliance on dense maps or voxel grids. In contrast, our pipeline works for sparse maps where the structure cannot be easily captured as compared to dense depth maps.

Interestingly, the trajectories generated by the RL based planner show a significant number of back and forth motions near regions of low continuity of visible mapped points. This is similar to manual

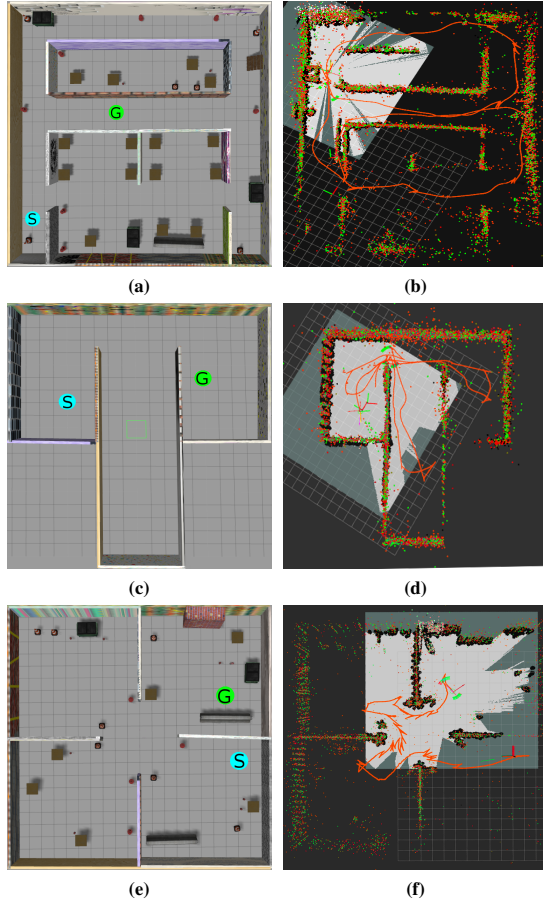


Figure 7: Results obtained with ORB SLAM using the Navigation Stack on maps 5 (a, b) 6 (c, d) and 7(e, f). The first column shows the environments. The robot was made to navigate from the start location (S) to the goal location (G). The second column shows the reconstruction from ORB SLAM and the robot’s estimated path after full exploration. motions employed with Monocular SLAM while using handheld cameras. This allows the SLAM algorithm to have continuity in the features tracked before moving into a new environment, thus leading to a better SLAM estimate.

The method is not only transferable to different maps but over different feature based Monocular SLAM algorithms as well. Values learned using one Monocular SLAM system are transferred to another, without any difficulties. In terms of learning SLAM failure safe behaviour, our methods showed about 37% improvement in comparison with a state-of-the-art Next-best-view planner and about 45% in comparison with Supervised Learning based planners. We also show the superiority of our approach over methods which use wheel odometry for multi-sensor fusion or relocalization.

In the future, we look forward to improving the performance of the system by integrating more recent techniques in Deep Reinforcement Learning and Deep Policy Gradients, which would allow us to add more parameters in the equation and explore into continuous space planning. Further work can also be carried out in studying how well these method scales to direct Monocular SLAM methods. Moreover, this method can be scaled to aerial robots and learn

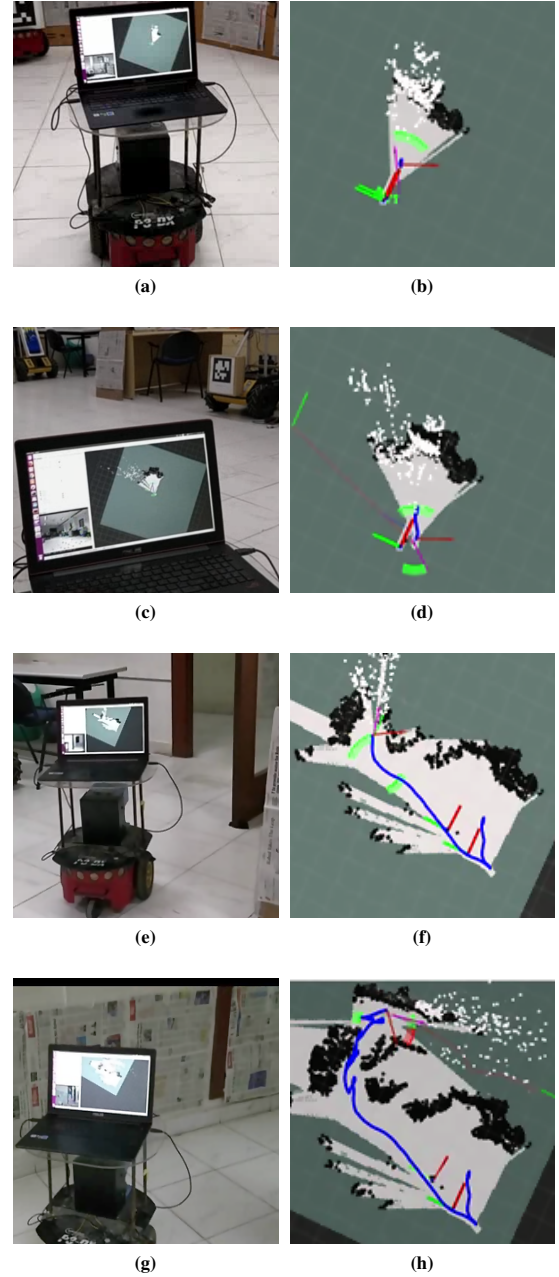


Figure 8: Results of real time exploratory run using the RL based Planner. First row shows the robot during its navigation from a large room to an open corridor. The second row shows trajectory and map estimates using ORB SLAM with RL based Planner where the predicted failures followed by back and forth motion around the opening of the room can be seen (as shown in (g) and (h) respectively).

SLAM failure safe behaviour for drones. In such cases, one could explicitly make use of the fact that better baselines between views can be obtained due to the horizontal/vertical motion capabilities of a drone as compared to just the forward/backward motion in case of ground robots.

REFERENCES

- [1] Pieter Abbeel and Andrew Y Ng. 2004. Apprenticeship learning via inverse reinforcement learning. In *International Conference on Machine Learning (ICML)*.
- [2] Benjamin Charrow, Gregory Kahn, Sachin Patil, Sikang Liu, Ken Goldberg, Pieter Abbeel, Nathan Michael, and Vijay Kumar. 2015. Information-Theoretic Planning with Trajectory Optimization for Dense 3D Mapping. In *Robotics: Science and Systems*.
- [3] Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning* (1995).
- [4] Gabriele Costante, Christian Forster, Jeffrey Delmerico, Paolo Valigi, and Davide Scaramuzza. 2016. Perception-aware Path Planning. *IEEE Transactions on Robotics* (2016).
- [5] Jakob Engel, Thomas Schöps, and Daniel Cremers. 2014. LSD-SLAM: Large-scale direct monocular SLAM. In *European Conference on Computer Vision*.
- [6] Jonathan D. Gammell, Siddhartha S. Srinivasa, and Timothy D. Barfoot. 2014. Informed RRT*: Optimal Incremental Path Planning Focused through an Admissible Ellipsoidal Heuristic. (2014).
- [7] Bharath Gopalakrishnan, Arun Kumar Singh, and K Madhava Krishna. 2014. Time scaled collision cone based trajectory optimization approach for reactive planning in dynamic environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- [8] Vadim Indelman, Luca Carlone, and Frank Dellaert. 2015. Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments. *The International Journal of Robotics Research* (2015).
- [9] Michael Kaess, Ananth Ranganathan, and Frank Dellaert. 2008. iSAM: Incremental smoothing and mapping. *IEEE Transactions on Robotics* (2008).
- [10] Georg Klein and David Murray. 2007. Parallel tracking and mapping for small AR workspaces. In *IEEE and ACM International Symposium on Mixed and Augmented Reality*.
- [11] Nathan Koenig and Andrew Howard. 2004. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- [12] Thomas Kollar and Nicholas Roy. 2008. Efficient Optimization of Information-Theoretic Exploration in SLAM. In *AAAI Conference on Artificial Intelligence*.
- [13] Steven M LaValle and James J Kuffner. 2001. Randomized kinodynamic planning. *The International Journal of Robotics Research* (2001).
- [14] Cindy Leung, Shoudong Huang, and Gamini Dissanayake. 2006. Active SLAM using model predictive control and attractor based exploration. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- [15] Cindy Leung, Shoudong Huang, Ngai Kwok, and Gamini Dissanayake. 2006. Planning under uncertainty using model predictive control for information gathering. *Robotics and Autonomous Systems* (2006).
- [16] Christian Mostegel, Andreas Wendel, and Horst Bischof. 2014. Active monocular localization: Towards autonomous monocular exploration for multirotor mavs. In *IEEE International Conference on Robotics and Automation*.
- [17] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. 2015. ORB-SLAM: a versatile and accurate monocular SLAM system. *IEEE Transactions on Robotics* (2015).
- [18] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* (2011).
- [19] Vignesh Prasad, Rishabh Jangir, Balaraman Ravindran, and K Madhava Krishna. 2017. Data Driven Strategies for Active Monocular SLAM using Inverse Reinforcement Learning. In *International Conference on Autonomous Agents and Multi-Agent Systems*.
- [20] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. 2009. ROS: An Open-source Robot Operating System. In *ICRA workshop on open source software*.
- [21] C. Roesmann, W. Feiten, T. Woesch, F. Hoffmann, and T. Bertram. 2012. Trajectory modification considering dynamic constraints of autonomous robots. In *The German Conference on Robotics, ROBOTIK*.
- [22] Richard S Sutton and Andrew G Barto. 1998. *Reinforcement learning: An introduction*. MIT press Cambridge.
- [23] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning* (1992).
- [24] Christopher John Cornish Hellaby Watkins. 1989. *Learning from delayed rewards*. Ph.D. Dissertation. University of Cambridge England.
- [25] Stephan Weiss, Davide Scaramuzza, and Roland Siegwart. 2011. Monocular SLAM-based navigation for autonomous micro helicopters in GPS-denied environments. *Journal of Field Robotics* (2011).
- [26] Georges Younes, Daniel Asmar, Elie Shammas, and John Zelek. 2017. Keyframe-based monocular SLAM: design, survey, and future directions. *Robotics and Autonomous Systems* (2017).