Active Perception for Rapid Deployment of Robot Teams in Unknown Environments

Jeffrey Delmerico[1], Elias Mueggler[1], Alessandro Giusti[2], Julia Nitsch[1],
Luca Maria Gambardella[1], and Davide Scaramuzza[2]

***Abstract*— We describe a system that utilizes rapid learning and active perception to plan a path for a ground robot through unknown terrain, using observations from a flying robot. In search and rescue missions, the time from arrival at the disaster site to the delivery of aid is critically important, so our approach focuses on minimizing this response time. Due to the unknown environments present in these scenarios, we propose a terrain classifier that can be trained and deployed *quickly*, based on data collected *on the spot*. We demonstrate that we can launch our aerial robot, gather data, train a classifier, and begin building a terrain map after only 60 seconds of flight. Our system also utilizes *active* exploration by the flying robot, where it maps the terrain and elevation in the environment in order to explicitly minimize the combined aerial exploration and ground robot transit time. The terrain class from our rapid classifier and the elevation estimates in the map are used to generate feasible and efficient paths for the ground robot. Our overall system is capable of being deployed rapidly in a previously unknown environment, with no prior knowledge or map.**

Supplementary Material

An accompanying video of the system is available at: `https://youtu.be/4cxDHPIEUx4`.

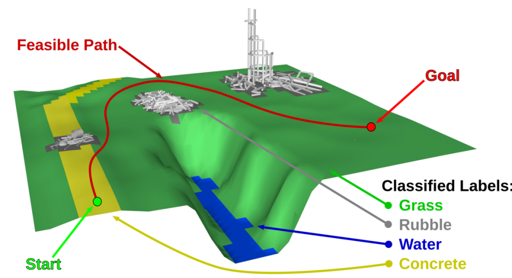

(a) Terrain Map

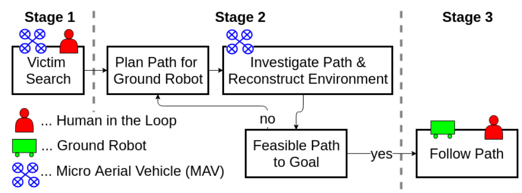

(b) System Overview

Fig. 1: Overview of the proposed system. Fig. 1a shows a diagram illustrating our intended terrain map output. It includes the elevation information, terrain classification, and the path found for the ground robot. Fig. 1b shows the workflow of the collaborative team.

I. Introduction

In disaster environments or search and rescue scenarios, time is a critical factor in the success of the first responders. One challenge is that the environment may have been altered by the disaster (e.g., an earthquake or a mudslide), potentially invalidating any prior maps. Consequently, robotic systems that can benefit first responders must be capable of gathering and using data *on demand, without reliance on a priori maps*.

The approach proposed here is motivated by the need to deliver a fast unmanned response in a previously unexplored environment using a collaborative robot team. Our system operates in three stages, shown in Fig. 1: an operator initially flies the MAV searching for a goal location for the ground robot (e.g. a victim) while collecting training data for the terrain classes present in the environment [1]; the MAV then explores the environment autonomously while building a map with elevation and terrain classes until it has found a feasible path for the ground robot; finally the ground robot executes the path to the goal. To minimize the overall response time from system deployment to delivery of aid, the exploration phase seeks to minimize both the estimated travel time for the ground vehicle and the time required to explore the map from above [2]. This allows us to adapt our classification to the terrain that is present in the search and rescue environment, without relying on a priori maps or classifiers.

II. Approach

Our robot team consists of a lightweight MAV and an all-terrain ground vehicle that can climb moderate grades and traverse small obstacles. Our MAV [3] is equipped with a downward-looking camera, and flies in autonomous and vision-assisted manual flight modes using the visual odometry pipeline SVO [4]. The images from this camera are additionally used for terrain classification, and for elevation mapping using the keyframe-based monocular dense reconstruction pipeline REMODE [5].

Our terrain map is a finite region of the ground surface, discretized into a 2D grid of uniformly sized cells. Our software interface allows the user to select a rectangular region of an image and label it as a terrain class, which is then associated to a region of the map due to the known pose of the MAV. All patches that subsequently project to the labeled cells are collected together as training data.

The authors are with [1]Robotics and Perception Group, University of Zurich, Switzerland (`http://rpg.ifi.uzh.ch`), and [2]Dalle Molle Institute for Artificial Intelligence (IDSIA), USI-SUPSI, Switzerland

This research was funded by the the Swiss National Science Foundation through the National Center of Competence in Research Robotics (NCCR).

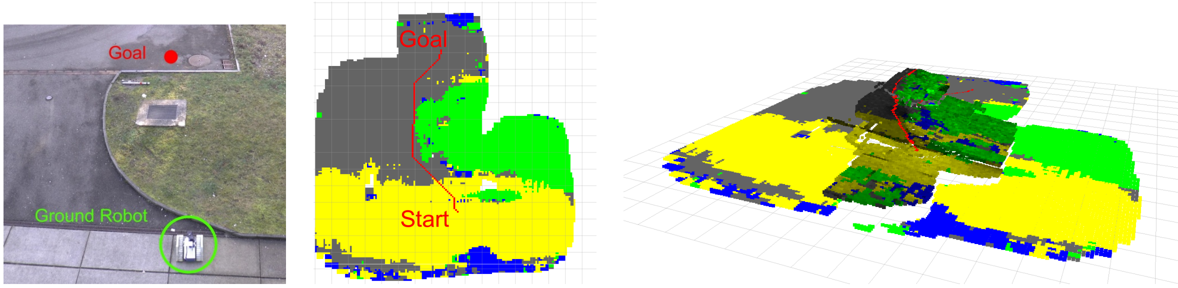

(a) Overview (goal in red) (b) Initial Classification (c) Elevation along path

Fig. 2: Driveway Experiment: While a direct path would be shorter, the chosen path (in red) would have lower response time due to the speed of the ground robot on different terrains. The terrain classes are *concrete* (yellow), *pavement* (gray), *grass* (green), and *water* (blue). Elevation mapping is only performed along the path for the ground robot, as visible in (c).

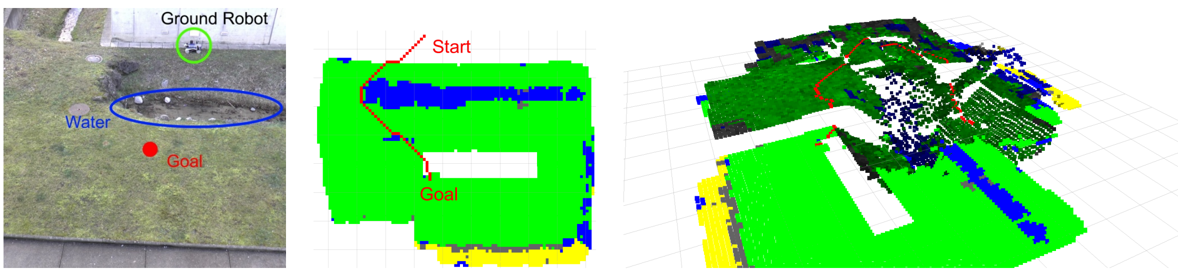

(a) Overview (goal in red) (b) Initial Classification (c) Final Result

Fig. 3: Canyon Experiment: The ground robot cannot cross the water. Our system finds a feasible path for the ground robot that stays on the grass. The terrain classes are *concrete* (yellow), *pavement* (gray), *grass* (green), and *water* (blue). The path is shown in red in (b) and (c). Elevation mapping is only performed along the path for the ground robot, as visible in (c).

After gathering training patches for the terrain classes, we train a *feature-based* classifier and apply it to randomly sampled patches from the image stream to classify the map cells that they project to.

We discretize the map into potential waypoints for the MAV to visit for 3D reconstruction, so the map is decomposed into a non-overlapping grid of patches. The problem of efficiently finding a feasible path for a ground robot is one of minimizing both the path traversal time, and the number of waypoints visited (and therefore the MAV flight time). We use both the estimated terrain class and elevation to determine traversable paths in the map, and estimate their costs in terms of response time. We propose an exploration strategy that utilizes a search over candidate paths in order to explicitly minimize the *total response time* of the system, not just the path cost.

III. Experiments

We successfully tested our system in two outdoor scenarios, to demonstrate the main capabilities of our system and verify that our path planner avoids untraversable terrain classes and handles significant elevation changes. We followed the procedure in Sec. II for both datasets, and were able to train the classifier in 60.44 seconds and 60.12 seconds, respectively. We then proceeded to survey the environment as we would in a mission scenario, and classify patches sampled from the image stream, while accumulating terrain class probability estimates in the map cells to which they project.

The first scenario (Fig. 2a) demonstrated the system's capability to distinguish different terrains: a straight line path through the grass might seem faster (and is feasible for the ground robot), it is better to stay on concrete since the ground robot can drive twice as fast on a hard surface. When driving the path with the ground robot, we found that the path on concrete was in fact 50% faster than the straight path over grass. The second scenario (Fig. 3a) included steep terrain and water, which cannot be traversed by the ground robot. The terrain classifier safely distinguished pavement, grass, concrete, and water, leading to a feasible path for the ground robot (Fig. 3b). The elevation mapper ensured that the path remained in regions that were flat enough for the ground robot to traverse (Fig. 3c), and we could successfully follow it with the ground robot.

References

[1] J. Delmerico, A. Giusti, E. Mueggler, L. Gambardella, and D. Scaramuzza, ““On-the-spot training” for terrain classification in autonomous air-ground collaborative teams,” in *ISER*, 2016.

[2] J. Delmerico, E. Mueggler, J. Nitsch, and D. Scaramuzza, “Active autonomous aerial exploration for ground robot path planning,” in *ICRA*, 2017.

[3] M. Faessler, F. Fontana, C. Forster, E. Mueggler, M. Pizzoli, and D. Scaramuzza, “Autonomous, vision-based flight and live dense 3D mapping with a quadrotor MAV,” *J. of Field Robotics*, pp. 1556–4967, 2015.

[4] C. Forster, M. Pizzoli, and D. Scaramuzza, “SVO: Fast semi-direct monocular visual odometry,” in *ICRA*, 2014.

[5] M. Pizzoli, C. Forster, and D. Scaramuzza, “REMODE: Probabilistic, monocular dense reconstruction in real time,” in *ICRA*, 2014.