Final newsletter for the EUROPA2 project

This is the final newsletter about the EU-funded research project EUROPA2 (European Robotic Pedestrian Assistant 2.0). In this newsletter, we begin with a brief summary of the project goals and previous status, before giving an overview of the whole system, and how we achieve the project goals using all the previously mention parts.

The Project

Recent years have witnessed an increased interest in the commercial use of robots operating in more complex and dynamic environments. The increased complexity of the environment introduces several challenges that need to be addressed before the robots can be safely deployed. Such challenges include representing the surrounding environment, the ability to localize in it and plan certain actions, like navigating to a goal. Moreover, both the semantic representation of the environment and the plan need to be robust enough to account for any deviations that occur due to structural changes in the scene or the presence of other agents.

The EUROPA2-project is concerned with several aspects of mobile robotics:

- Online semantic scene interpretation for more efficient navigation.
- New means to combining publicly available prior map models with the perceptions of the robot, enabling the system to maintain and exploit map knowledge without having mapped a city beforehand.
- Bootstrapping the navigation system with publicly available maps to avoid requiring to map the environment with the robot in advance.
- Life-Long operation of the whole robotic system, including map management, navigation, and calibration.
- Perception under different weather conditions and seasons.
- Introspective classification for robust decision making, i.e., considering the certainty of classification results appropriately.
- To extend and improve existing online methods for detecting and tracking dynamic objects towards close-range and partially visible pedestrians and other traffic participants.
- Develop the foundations for representations of dynamic urban environments for autonomous robots. This includes better models for the interaction between people and the robot as well as with other traffic participants such as cars.

In the previous newsletter, we presented some of the progress made in terms of semantic scene interpretation, intelligent energy management, life-long self-calibration, and learning interaction models for navigation.

In this newsletter, we report on how the different system parts are integrated together.

Progress

Semantic Mapping and Segmentation

In order for the robot to be able to interact with the surrounding environment, it needs to be able to interpret the scene around it. This includes the ability to identify the main road, sidewalks, flanking facades, parked cars, traffic signs, etc. We developed a real-time algorithm that relies on the stereo camera of the robot to be able to recover basis scene analysis. We use a deep learning approach for monocular semantic segmentation that is both accurate and online. Below are example images from running the semantic segmentation module live on the EUROPA2 robot.

Safety is a major concern in the operation of autonomous mobile robots. In the context of the EUROPA2 project, the robot will be operating in an urban environment which is characterized by the presence of pedestrians. As such it is crucial for the robot to be able to detect, and track people so as to avoid collisions during motion. We developed a robust detector with a strong focus on efficiency. The detector is able to process a video feed in real time. Below are some example images of running the detector. Rectangles highlight the detected pedestrians.

Not only is it important to detect the presence of pedestrians, we must also be able to track their motion in order to be able to plan the robot's motion safely around them. We developed a pedestrian tracker, that fuses 3D-laser measurements from the robot with the output from the pedestrian detection. The information from the tracker is provided for the motion planning modules.

To ensure safe operation with cities, the robot must be able to both detect and recognize the different traffic signs in order to act upon them. For this purpose, we developed a deep learning framework for traffic sign detection and recognition. The approach is able to run online and in parallel with the semantic segmentation and pedestrian tracking modules.

Motion Planning

The planning system on the robot is divided into three modules; a global planner, and intermediate planner, and a low-level planner. The global planner computes a trajectory on the GeoAutomation and OpenStreetMap maps. It does not account for obstacles along the planned trajectory, rather treats the map as a graph structure and attempts to find a path from the starting point to the goal destination. Avoiding obstacles is handeled by the intermediate planning module. It uses the traversability analysis information to create feasible paths and sends target points to the low-level planner. Finally, the low-level planner is responsible for generating motion commands, while taking into account dynamic objects and pedestrian interactions in the surrounding neighborhood.

In addition to the global planner, we develop a local planner, called Social Navigation Planning. The social navigation planning uses the information from the pedestrian tracker to predict future behavior of pederstrians by reasoning about their interactions among each other and with the robot. Similar to the global planner, it is part of a three-layer motion planning framework.

As the social navigation planning is interaction aware, it therefore requires motion models for the surrounding pedestrians. These models are trained by real-life demonstrations captured using only the robot's onboard sensors. Below is an image displaying the interactive motion planning module. The planned path for the robot is visualized in orange, the prediction for the other pedestrian is shown in blue.

Localization and Navigation

Fitting all the localization and navigation components together, we developed a 3D localization module that localizes the robot on GeoAutomation maps, self-captured 3D maps and a combination of the two. Furthermore, to be able to avoid obstacles in the robot's path, we developed a traversability analysis module that uses 2D scanners. It has low CPU requirements during operation (approximately 25%), and is fully integrated into the platform providing input for the planning modules.

In order to improve on the localization capabilities of the system in existing maps, it is advantageous to remove dynamic/moving objects from the sensor data. For this purpose, we developed a module to detect moving objects in 3D scans. An Example 3D scan from the robot, with four people walking behind it marked as being dynamic (yellow) is shown below.

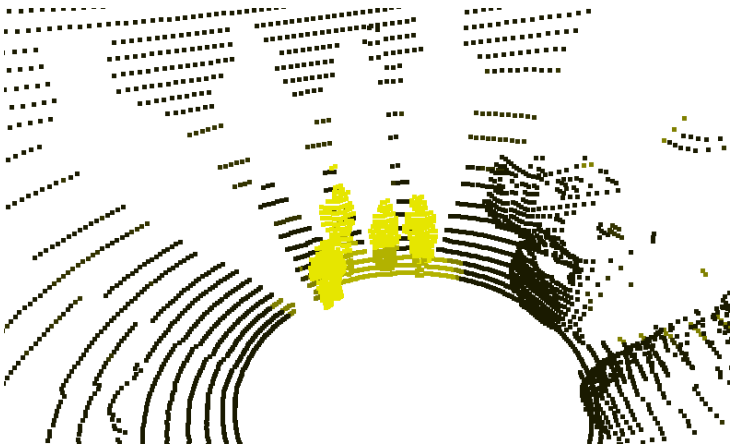

One important prerequisite to perform autonomous navigation is to have a map of the environment the robot should operate in. Since acquiring maps with a mobile robot comes with a large effort at city scale, we developed a module that uses the pre-existing road network from OpenStreetMaps to build an initial map for the robot. The map stores semantic information regarding the nature of the road like the type of the road, the presence of a side-walk, locations of zebra crossings, and so on.

For the robot to be able to localize in the given map, we implemented a localization method based on Monte-Carlo localization that matches the robot's 3D sensor data against a sparse point cloud from the GeoAutomation map. The image below depicts the EUROPA2 robot localizing in the GeoAutomation maps. One can clearly see the estimate of the robot location (red), the GeoAutomation tracks (black) and the current sensor measurements (blue).

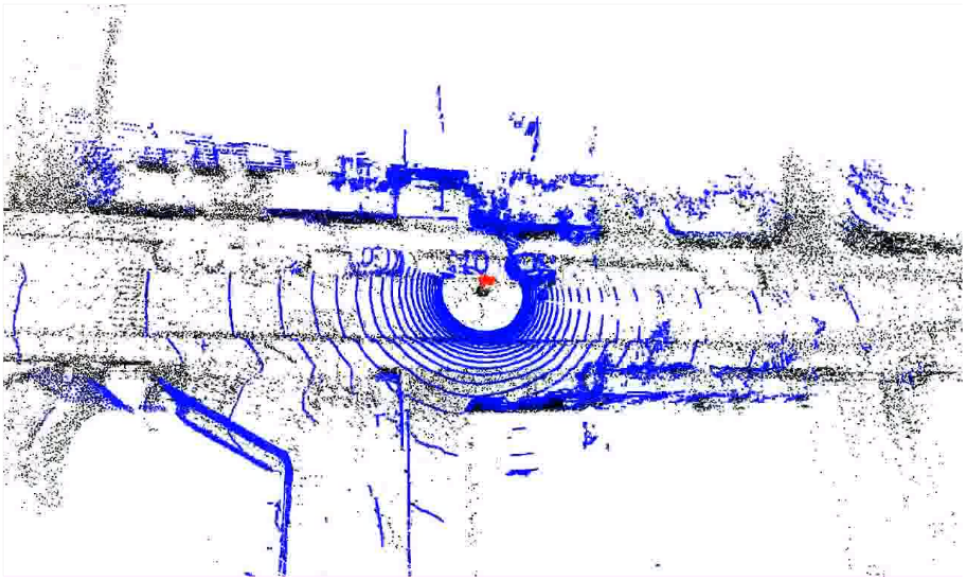

One important key capability to enhance or extend publicly available maps is to estimate a consistent model of the robots sensor observations and relate those information to the underlying map. This gives us the capability to adjust the map with new information in case of changes, and adding more information where information was lacking, for instance inside buildings. We developed a global optimization approach that is able to build 3D maps of the environment using the 3D laser information, GPS measurements and IMU measurements from the robot's sensors. We modify the OpenStreetMap road network by adding links relating the robot's traversal. This enables us to construct consistent models from the robot's sensor data, while increasing the accuracy of the information. Below a figure illustrating our map structure before (top left) and after extending it (bottom) with data collected from our robot platform (top right). The extended map includes the 3D information collected by the robot and includes new edges in the road network based on the driven trajectory of the robot.

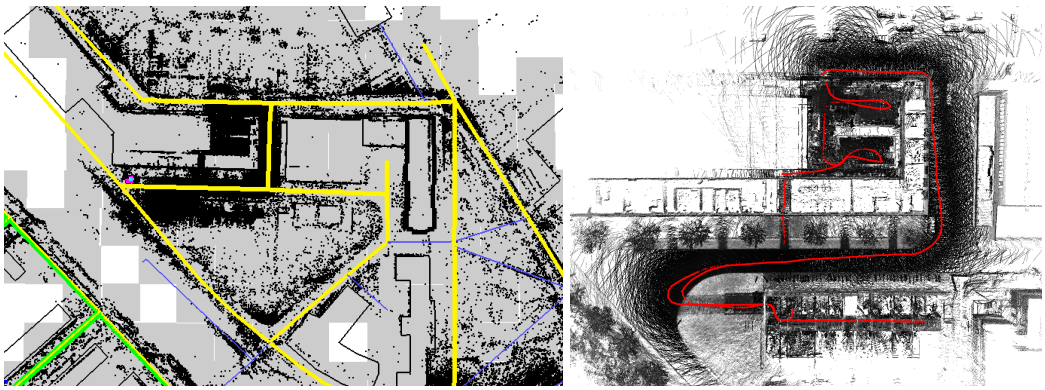

Thank you for your interest in the EUROPA2-project.