Second newsletter for the EUROPA2 project

This newsletter is about the EU-funded research project EUROPA2 (European Robotic Pedestrian Assistant 2.0). In this second newsletter, we begin with a brief summary of the project goals and previous status, before introducing current progress.

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 using publicly available maps for both localization and planning, developing a navigation system that relies mainly on vision, and supports life-long operation in urban environments. Finally, we presented a system for people detection that is able to both accurately and efficiently detect pedestrians through real-time video feed.

In this newsletter, we report on further progress made during the previous development months.

Progress

Semantic Scene Interpretation

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. This task exploits location priors based on map information, and using active perception to be able to run in an online scenario. Below is an example of the semantic scene interpretation where traffic lights are detected along with the signal displayed.

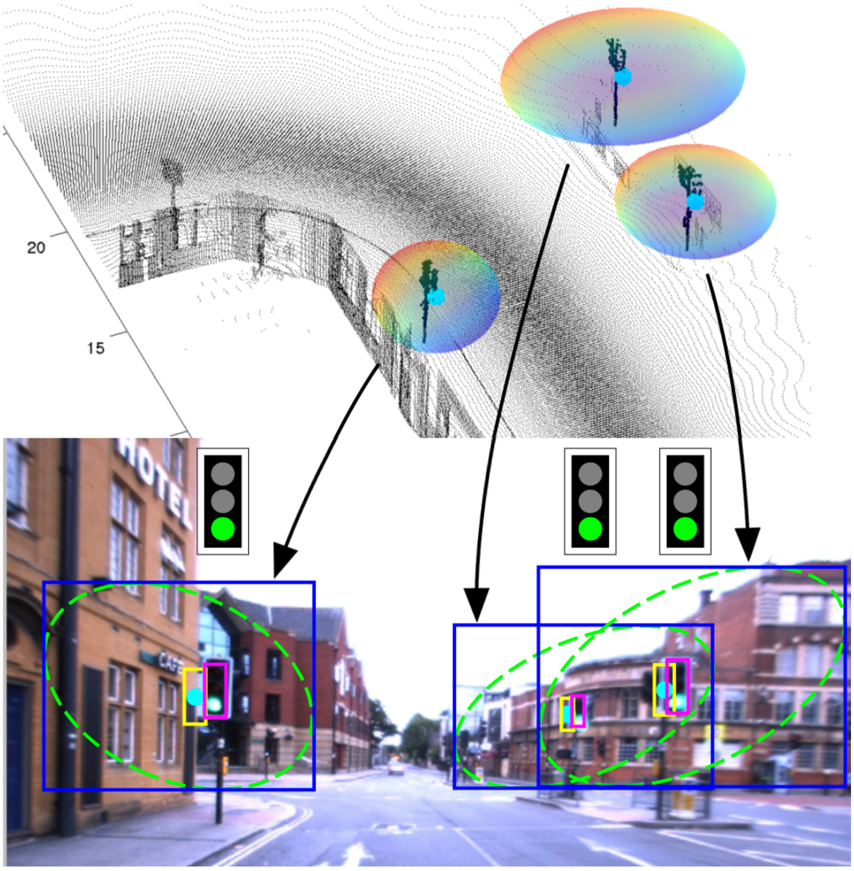

Intelligent Energy Management

Energy consumption is a major constraint in the deployment of mobile robots, especially for long-term tasks. We implemented a new framework for estimating the energy consumption using only information from publicly available maps. The developed approach uses density regression to model the energy consumption of the robot on a set of planned trajectories. The goal is to maximize the battery time. The developed energy consumption measurements are highly variable, due to external factors such as the time of the day, pedestrian traffic, and weather conditions. A particularly crowded environment, for example, will require the robot to stop more often to avoid people, increasing the energy consumption.

Learning Interaction Models for Navigation

The robot needs to be behave in a socially accepted manner, as it will be operating in urban environments which have the tendency of being crowded with either pedestrians, cyclists or cars. To address this issue, we implemented a local planning module that takes into account the movements of pedestrians and adjusts its plan according to an interaction model for pedestrians. The image below visualizes the sensor data in a situation with three pedestrians. The red arrows show the current movement directions. The robot plans a path according to the assumed green trajectories.

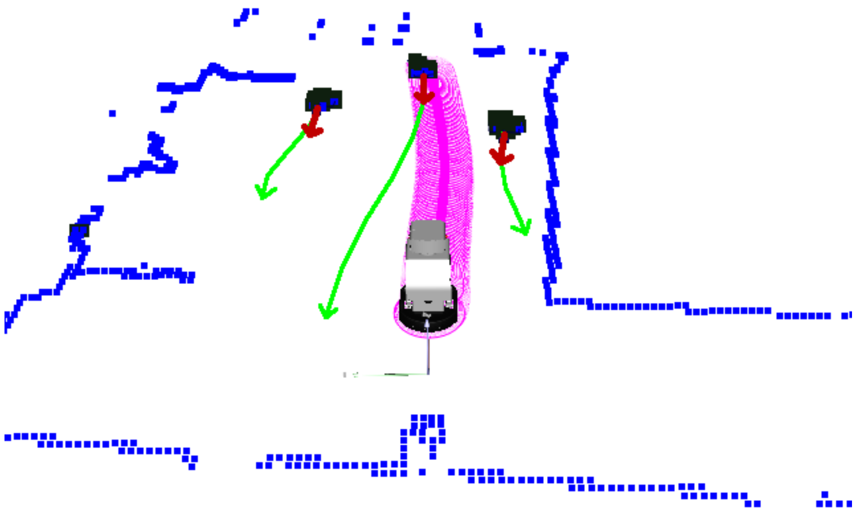

The probabilistic model that is used to predict the behavior of interacting pedestrians is a feature-based maximum entropy model. Features are used to capture characteristics of trajectories and interactions between different agents. Training the model according to the principles of maximum entropy theory helps us to reduce the generalization error of the model. The model is trained with data recorded with the onboard sensors of the EUROPA2 platform without having a discrete set of possible target locations for pedestrians. This gives the approach the desired robustness for real-time applications and allows for a fast deployment in previously unseen areas.

To accomplish the task of safely navigating in urban environments, we integrate the information from other modules such as pedestrian tracks, and semantic labels. Available information is fused to produce safe motion plans that enable us to transition from simple rule-based planning, e.g., how to cross a street, to interaction-reliant planning, e.g., how to navigate through crowds.

Life-Long Self-Calibration

One major focus of the EUROPA2 project is life-long operation. This mandates the need of the robot to self-calibrate and detect failures during operation. There are several challenges which we tackle to solve the self-calibration problem. First is the problem of unknown time delays between different sensor readings. Another problem is the observability of the state space. If the state or calibration space is not entirely observable, the information available to correctly calibrate the sensors is not sufficient. Finally, the problem of sensor data associations, which refers to associating data across sensors. This is not a major problem in in-laboratory calibration scenarios using specific markers in the environment. However in an unmodified operation environment, we need to be able to properly associate the different sensor readings. We solve these challenges using a mixture of mutual information maximization concept, with Iterative Closest Point (ICP) based point-to-point associations, and Binary Robust Invariant Scalable Keypoints (BRISK) features.

CNN based detection, classification and semantic analysis

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, It is crucial for the robot to detect them, but also the road structure, obstacles, and other objects may need to be detected without having priors on the local map structure.

In the previous period we focussed on the development of a robust detection mechanism for pedestrians. We developed a robust detector with a strong focus on efficiency, to run real-time on CPU. Also traffic sign detection was investigated in this respect. While the performance was satisfactory, it was not easy to increase the performance to avoid false positives or negatives.

The recent uprising of convolutional neural networks, who have proven to be particularly powerful for classification or segmentaion, given sufficient amount of data for training, has encouraged us the create novel detectors based on this paradigm. At the moment the first versions of several classifiers have been integrated

- pedestrian detector, using fixed+CNN based filters, and an Adaboost approach for detection,
- a traffic sign detector, using 9 super classes for detection, and a subsequent cnn-based recognition step for identification,
- a semantic analyser, trained for road area, foreground objects (including pedestrians), buildings, trees, and sky.

The next picture gives a snapshot of the robot operating live on a test sequence where a few persons are holding traffic signs, and moving around the university complex. Bottom right is the robot itself, the central screen in the middle operates the software and visualises the output of the detectors. On Top left is the pedestrian detector, on the top right the semantic analysis, and on the bottom left the traffic sign detection and recognition. The development and integration process will be continued and tested in even more realitistic environments.

Thank you for your interest in the EUROPA2-project. In the next and final newsletter, we will provide information about more aspects of the project and summarize the final outcomes of the project.