CS 8803 Machine Learning Applications for Robotics

Credits	3 Lectures hours 0 Lab hours	Semester Credit Hours:	3
Faculty	Prof. Cédric Pradalier		
Bibliography	Bishop, Christopher M. *Pattern recognition and machine learning*. Vol. 4, no. 4. New York: Springer, 2006. Thrun, Sebastian, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. Vol. 1. Cambridge, MA: MIT press, 2005.		
Assessment	2 written exams: 2 x 10%, Students start with 80% of project credit. Penalty for not returning a homework -30%, for not returning a project: -50%.		
Course Objectives:	This class will present computational tools, mostly from the machine learning toolbox, useful for the development of robotic applications. The objective is to get a hands-on experience with these tools through some theoretic background and a lot of project work.		
Pre-requisites:	This is a CS class, hence it assumes a familiarity with programming and computers, as well as a reasonable background in maths, in particular linear algebra and geometry. Support will always be available for technical issues. Linear Algebra: being at ease with manipulating matrices Programming: the projects will require a significant amount of programming in C++ and Python. The students are expected to either know these languages or have the motivation to learn them on the fly. Environment: all the projects will be run under Linux. The students are expected to either know this OS or have the motivation to learn it on the fly.		
Weekly organization:	Outside of the project work, the class will start with a theoretical session (80-120min) and then head to the computer room to start on the week mini-project. The class will include 3 larger projects, each of them done over a 3 week timespan. In this case, the full duration of the class will be dedicated to project work in the computer room with faculty support and interventions on the white board.		
Infrastructure:	The (mini-)projects will be conducted in groups of 2 students in a simulation environment installed in GTL's computer room. The simulation will run on top of V-Rep (www.v-rep.eu) and will use ROS (www.ros.org) to communicate with the simulated robots. For those highly motivated, work on real robots can probably be arranged. Please make yourself known early in the semester.		

Description and Class Schedule:

Class Nr	Week	Topic	Homework and Projects
1	34	Introduction	Introduction to the Development Environment
2	35	Linear Regression	Plane fitting using RGB-D Sensors
3	36	System Identification and Control	Robotics System Identification
4	37	Model Finding: RANSAC/Hough	Object finding using RGB-D sensors
5	38	Mid-Term 1	Introduction to Project 1
6	39	Project 1 + Bayesian Filtering	Project 1
7	40	Project 1	
8	41	Bayesian Filtering for Localisation and Mapping: Kalman and Particle Filter	Robot Localisation using Kalman Filter
9	42	Mid-term 2	Robot Localisation using Particle Filter
	43	RECESS	
10	44	Introduction to Convolutional Neural Networks (CNN)	Introduction to Project 2
11	45	CNN for Control / Regression	Project 2
12	46	Project 2	
13	47	Unsupervised Learning: From Clustering to EM	Introduction to Project 3
14	48	Project 3	Project 3
15	49	Project 3	
	50	Finals	

Project description:

Project 1: Probabilistic Mapping. Considering a robot with known perfect localisation and a RGB-D sensor (Kinect), use Bayesian filtering (BF) to estimate a digital elevation map (BF for a Gaussian continuous variable) and traversability (BF for a binary variable) using the tools presented in weeks 34 to 37.

Project 2: Using Convolutional Neural Network. The goal of this project will be to use off-the-shelf CNN libraries (Caffe) to train and use a CNN. There will be two sub-components of the project. In the first stage, a CNN will be trained to perform a classification task to detect the drivable area of an environment from camera. In a second stage, a CNN will be trained to learn a navigation behaviour by imitation. In both case, the problem will be framed as a classification task.

Project 3: Integration and Treasure Hunt. The objectives of this final project is to integrate some of the components learnt until here into a complete application design to search an environment for hidden objects (this could also be seen as a demining application). The simulated robot will be a large truck with a 5 DoF robotic arm. Autonomous Navigation and Object Detection will have to be implemented by the students.