Syllabus

Overview

Course title: Self-Driving Cars with Duckietown
Course ID: ETHx+DT-01x+1T2021
Method: Self-paced, Verified and audit tracks
Availability: 29 Nov. 2022 - 25. Nov. 2023
Course length: 9 modules plus introduction
Expected weekly effort: simulation track: 3-5 hours, hardware track: 6-10 hours
Supported hardware: Duckiebot with Jetson Nano 4GB (DB-J4) and Duckietown Navigation Starter pack (MOOC starter kit)

Intended Learning Outcomes

This course is designed to be hands-on. This means that you will actually have to *do* things in the real world with the information we will provide you. In the context of robotics, the "real" world could mean a simulated environment as well as the physical world.

By the end of this course, you will learn about the fundamentals of robotics. What qualifies as a robot? What do they have in common? What are the differences and implications of logical and physical autonomy architectures? How significant is the difference between what *should* happen in theory and what *actually* happens, and how do we go from understanding to implementation?

Self-driving Cars with Duckietown is a "grand tour" of robotics. The focus is on breadth rather than depth. We highlight how robotics is a system-level discipline where many components interact with each other to produce real-world autonomous behaviors. Any of these "components" could absorb an entire professional career.

Finally, we want to provide an understanding of the difference in approaches to solving the challenges of autonomy. The buzz surrounding AI is high. What is its place in robotics? How does using reinforcement or imitation learning compare with the "traditional" robotics approaches to given tasks?

Robotics and AI are fast-paced fields of research and this course is just the beginning of a learning adventure.

Course Content

MODULE 0: Welcome to the course

- Course introduction by Prof. Emilio Frazzoli
- You will familiarize with the logistics and navigation interface of the course resources
- You will start a learning journey in the world of robot autonomy with Duckietown

MODULE 1: Introduction to self-driving cars

- The potentials and challenges
- Levels of autonomy
- The vision for autonomous vehicles (AVs)
- Activities: You will set up your learning environment and duckiebot and make your first submission to the AI Driving Olympics

MODULE 2: Towards autonomy

- Making a robot
- Sensorimotor architectures
- Stateful architectures
- Logical and physical architectures
- Application: You will create your own Braitenberg agent for the task of avoiding duckies and evaluate how your agent compares to other submissions

MODULE 3: Modeling and Control
- Introduction to control systems
- Representations and models
- PID control
- Application: You will design an odometry function and PID controller to command your Duckiebot's heading and lateral positions

MODULE 4: Robot Vision

- Introduction to projective geometry
- Camera modeling and calibration
- Image processing

- Application: You will explore image processing techniques to implement visual lane servoing - controlling your Duckiebot to drive within markings (Lane Following)

MODULE 5: Object Detection

- Introduction to Neural networks
- Convolutional neural networks
- One and two stage object detection
- Application: You will train a convolutional neural network (CNN) to detect duckies and integrate your model with ROS to run onboard your Duckiebot

MODULE 6: State Estimation and Localization

- Bayes filtering framework
- Parameterized methods (Kalman filter)
- Sampling-based methods (Particle and histogram filter)
- Application: You will compare state estimation algorithms combining the dynamics and sensor data of your Duckiebot in order to predict its pose as it travel through the world

MODULE 7: Planning I

- Planning problem formalization
- Graphs
- Application: You will create a collision checker to determine if your Duckiebot is crashing into an obstacle

MODULE 8: Planning II

- Probabilistic roadmaps
- sampling-based planning
- Application: You will tackle a variety of path planning challenges and leverage all the skills you've built thus far to navigate your Duckiebot in variety of simulated environments

MODULE 9: Learning by Reinforcement

- Markov decision processes
- Value functions
- Policy gradients
- Domain randomization

- Application: You will explore the capabilities and limitations of reinforcement learning models when applied to real-world robotics tasks such as Lane Following

Prerequisites

Basic Linux, Python, Git. We are going to:

- use terminal interface, so basic knowledge of Bash is required
- write "autonomy" code in Python
- pull Git repositories, fork, push, branch, etc.

Elements of linear algebra, probability, calculus. We are going to:

- use matrices to represent coordinate systems
- use notions of probability (marginalization, Bayes theorem) to derive perception algorithms for the Duckiebot
- write down equations of motion, which involve time ODEs (recognizing the acronym is a good start!)

Computer with native Ubuntu installation

- We are going to use Ubuntu 22.04 with a native (e.g., dual boot) installation
- Computer hardware requirements:
    - Minimum:
        - CPU: quad-core at 1.8 GHz
        - RAM: 4GB
        - Disk: 60GB
        - GPU: compatible with OpenGL 2.1+
    - Recommended:
        - CPU: quad-core at 2.2 GHz
        - RAM: 8GB
        - Disk: 120GB
        - GPU: compatible with OpenGL 2.1+

The Staff

Throughout the course you will meet instructors and staff from all the organizing institutions. Before starting, we wanted to very briefly introduce ourselves. Check out our informal introductions on Vimeo! In order of appearance:

- Prof. Emilio Frazzoli, ETHZ
- Dr. Andrea Censi, ETHZ (introduction, LinkedIn)
- Dr. Jacopo Tani, ETHZ (introduction, LinkedIn)
- Prof. Matthew Walter, TTIC (LinkedIn)
- Prof. Liam Paull, UdM (introduction, LinkedIn)
- Dr. Andrea Daniele, Duckietown (introduction, LinkedIn)

The Duckiebot and Duckietown

The physical Duckiebot and Duckietown are not required to follow and complete this course. However, following along with a real robot will help you learn much more. Here you can get a Duckietown MOOC Founder Edition Kit (robot + city track).

Grading policy

This course is pass/fail. A pass result is necessary to obtain a certificate and is achieved with a final grade of at least 60/100.

The final grade will be calculated as the average of the grades of the homework assignments. Each assignment will have the same weight towards the final grade.

How to get help

We will be answering questions related to

- *learning materials* on the EdX Discussions forum, and

- *technical questions* on the Duckietown Stack Overflow space. This is a private space and will require an invitation to join. We will process invitations in batches, periodically. If you really can't wait to receive an invitation you are welcome to join the Duckietown Slack and ask to have one sent in the #help-accounts channel.

Each module will be led by an instructor. The lead instructor will be the go-to person for addressing questions related to that module.

Additionally, Dr. Andrea Daniele will be responsible for supporting software challenges, and Dr. Jacopo Tani for the hardware-related ones. Please tag Jacopo and Andrea in your Stack Overflow question to expedite feedback.

To provide effective support, we will focus on answering questions posted in the above avenues, and ignore other requests for support (especially those received by email or direct messages on Slack).

Additionally, when using any forum we expect all learners to maintain a constructive, positive and professional attitude (but clearly duckie puns are always welcome!). We will ban / report / take all appropriate actions to maintain a high-quality learning environment.

Course handouts and operation manuals will be available in the grand Duckietown library. Links will be provided throughout the course inside the module units.