Welcome to CS 3630!

Fall 2022:

- *Online Book!*
- *Revamped projects!*

Course Instructors

Frank Dellaert, Professor

School of Interactive Computing

Stints at Skydio, Facebook, Google, now CTO at Verdant Robotics

Book co-author (and frequent 3630 teacher): Seth Hutchinson, Professor

School of Interactive Computing

Director of Institute of Robotics and Intelligent Machines (iRIM)

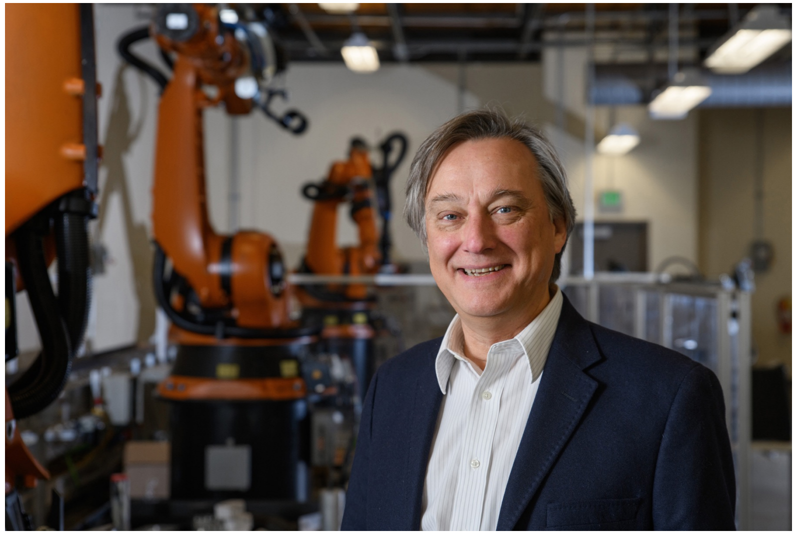

TAs: *Matthew King-Smith*, Isabella Poage, Srihari Subramanian, Abhineet Jain, Allison Fister, Aswin Prakash, Nitin Vegesna, Asha Gutlapalli, John Yi, Meher Nigam, Avinash Prabhu, Adwait Deshpande, Vivek Mallampati, Priyal Chhatrapati, Prannoy Jada

Robots are useful!

- Manufacturing
- Logistics (inventory, warehouse logistics, packaging)
- Transportation (self-driving cars)
- Consumer and professional services (cleaning, mowing)
- Health, independence and quality of life (exoskeletons, semi-autonomous wheelchairs)
- Agriculture

Robot Taxonomy

- **Industrial Robots**
- Service Robots
- Field Robots
- Humanoid Robots
- Medical Robots
- Self-Driving Cars
- Aerial Vehicles

http://www.kuka.com

Robot Taxonomy

- Industrial Robots
- **Service Robots**
- Field Robots
- Humanoid Robots
- Medical Robots
- Self-Driving Cars
- Aerial Vehicles

Robot Taxonomy

- Industrial Robots
- Service Robots
- **Field Robots**
- Humanoid Robots
- Medical Robots
- Self-Driving Cars
- Aerial Vehicles

http://www.frc.ri.cmu.edu/robots/

Robot Taxonomy

- Industrial Robots
- Service Robots
- Field Robots
- **Humanoid Robots**
- Medical Robots
- Self-Driving Cars
- Aerial Vehicles


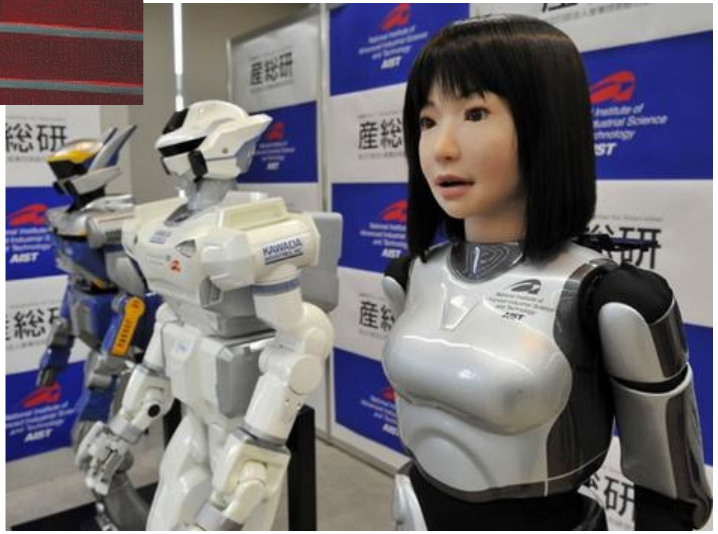

Robot Taxonomy

- Industrial Robots
- Service Robots
- Field Robots
- Humanoid Robots
- **Medical Robots**
- Self-Driving Cars
- Aerial Vehicles

Robot Taxonomy

- Industrial Robots
- Service Robots
- Field Robots
- Humanoid Robots
- Medical Robots
- **Self-Driving Cars**
- Aerial Vehicles


Robot Taxonomy

- Industrial Robots
- Service Robots
- Field Robots
- Humanoid Robots
- Medical Robots
- Self-Driving Cars
- **Aerial Vehicles**

In this class…

- We will not deal specifically with each of these robots.
- However, the mathematical and computational tools that we introduce can be applied to all the above robots.
- Likewise, the sensors and sensing methods that we will introduce can be applied to all the problems described above.

Class Website

Class Website

https://dellaert.github.io/22F-3630/

- Mostly externally facing
- Communication in class is via
  - Piazza (Q&A)
  - Canvas (esp. grades)
  - Gradescope

The course web site will always have the most up-to-date version of the schedule (next!).

Introduction to Perception and Robotics

Georgia Tech CS 3630 Fall 2022 edition

Home	Book	Resources
Syllabus	Schedule	Projects

Instructor: Frank Dellaert in Interactive Computing

TAs: Matthew King-Smith, Isabella Poage, Srihari Subramanian, Abhineet Jain, Allison Fister, Aswin Prakash, Nitin Vegesna, Asha Gutlapalli, John Yi, Meher Nigam, Avinash Prabhu, Adwait Deshpande, Vivek Mallampati, Priyal Chhatrapati, Prannoy Jada

Welcome to the homepage of CS3630, Fall 2022!

Course Description

This course covers fundamental problems and leading solutions to autonomous robot navigation – what and how must a robot perceive the world, and how can it use that information to navigate effectively.

Maintained by Frank Dellaert and the TAs of CS 3630

Based on a theme by orderedlist

Schedule

Class Schedule

There will almost certainly be changes to precise lecture topics, but quiz and project dates are unlikely to change.

The course web site will always have the most up-to-date version of the schedule.

Week	Month	Date	Day	Module	Topic	Slides	Reading	Quizzes	Projects
1	Aug	23	Tue		Introduction	-	-		Intro to Python
1	Aug	25	Thu	Sorting	State, probability, and actions	-	-		
2	Aug	30	Tue	Sorting	Sensor Models	-	-		P1: Probability and Decision Theory
2	Sep	1	Thu	Sorting	Bayes Law	-	-		
3	Sep	6	Tue	Sorting	Decision Theory	-	-		
3	Sep	8	Thu	Vacuum	Probabilistic Actions	-	-	Quiz 1	
4	Sep	13	Tue	Vacuum	Inference in HMMs	-	-		P2: HMM and MDP
4	Sep	15	Thu	Vacuum	Inference, MDPs	-	-		
5	Sep	20	Tue	Vacuum	MDPs	-	-		
5	Sep	22	Thu	Logistics	Omnidirectional wheels	-	-	Quiz 2	
6	Sep	27	Tue	Logistics	LIDAR sensing	-	-		P3: Particle Filters
6	Sep	29	Thu	Logistics	Monte Carlo Localization	-	-		
7	Oct	4	Tue	Logistics	Planning	-	-		
7	Oct	6	Thu	Duckiebot	Differential Drive	-	-	Quiz 3	
8	Oct	11	Tue	Duckiebot	Cameras and Image Processing	-	-		P3 due
8	Oct	13	Thu	Duckiebot	Computer Vision 101	-	-		
9	Oct	18	Tue	Fall Break	Fall Break	-	-		
9	Oct	20	Thu	Duckiebot	Inference with Deep Nets	-	-		
10	Oct	25	Tue	Duckiebot	Deep Learning	-	-		P4: Deep Learning
10	Oct	27	Thu	Vehicles	Autonomous Vehicles and SE(2)	-	-	Quiz 4	
11	Nov	1	Tue	Vehicles	Ackerman steering	-	-		
11	Nov	3	Thu	Vehicles	SLAM with 3D LIDAR	-	-		
12	Nov	8	Tue	Vehicles	Planning for Avs	-	-		P5: Pose SLAM and LIDAR
12	Nov	10	Thu	Drone	SE(3)	-	-	Quiz 5	
13	Nov	15	Tue	Drone	RRTs for motion planning	-	-		
13	Nov	17	Thu	Drone	Visual Odometry	-	-		
14	Nov	22	Tue	Drone	Trajectory Optimization	-	-		P6: RRT and/or Trajectory Optimization
14	Nov	24	Thu	Thanksgiving	Thanksgiving	-	-		
15	Nov	29	Tue	Extra	Guest Lecture TBD	-	-	Quiz 6	
15	Dec	1	Thu	Extra	Guest Lecture TBD	-	-		
16	Dec	6	Tue	Extra	Guest Lecture TBD	-	-		P6 due, all quizzes graded!

Six Modules

The class is organized into six modules, each of which focuses on a specific robot performing a specific application:

- A Trash Sorting Robot
- A Vacuum Cleaning Robot
- A Robot for Logistics (e.g., warehouse operations)
- The Duckiebot (a simple wheeled mobile robot)
- Autonomous Cars
- Drones

Trash Sorting Robot

- Pieces of trash arrive to a robot work cell.
- The robot’s job is to classify each piece of trash and move it to an appropriate bin.
- Simple, deterministic, high-level actions.
- Simple sensors.
- Uncertainty in sensor readings introduces probability into perception.
- Planning is the problem of choosing actions to minimize average costs over a long horizon.

Vacuum Cleaning Robot

- Robot actions are uncertain: command to move to living room might take the robot to the kitchen.
- Sensing is uncertain.
- Perception is addressed using Hidden Markov Models (HMMs).
- Planning is addressed using Markov Decision Processes (MDPs).

Simple Logistics Robot

- Mobile robot platform that moves in a warehouse.
- Omnidirectional wheels, so the robot can move in any direction at any time (not like a car).
- LIDAR sensing (includes uncertainty)
- Monte Carlo localization to deal with uncertainty when determining the robot's location.
- Planning is not so difficult, because warehouses are fairly regular and well organized.

Duckiebots

- The Duckiebot has two wheels (differential drive), and can only move “forward/backward” (can change the direction of “forward/ backward” by rotating).
- Rotation complicates the geometry of motion. We’ll introduce the appropriate mathematics to deal with rotations.
- For perception, we’ll introduce deep learning methods (very casual and superficial – not a deep dive into deep learning).
- In a post-pandemic world, with smaller lecture sizes, we hope one day to use the Duckiebots in the lab portion of this course.

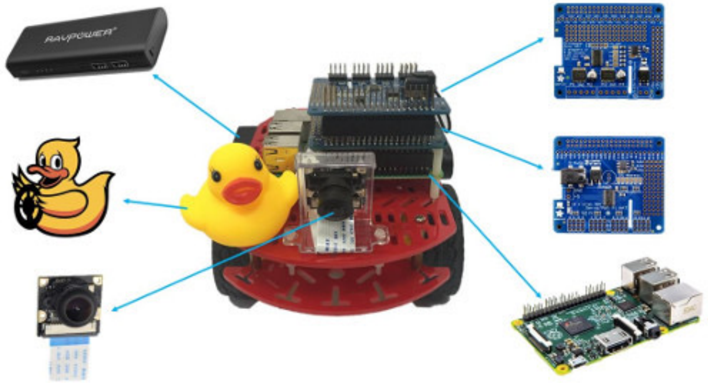

Autonomous Cars

- The robot is a fully instrumented autonomous (aka self-driving) car).
- The car has a non-zero turning radius (unlike the Duckiebot), and cannot move sideways.
- We'll use LIDAR sensors to determine the world state.
- Project will involve a large, real-world data set used for autonomous driving research.

Drones

- Drones fly in 3D.
- Dealing with rotations in 3D is tricky, but we'll introduce the math to deal with this.
- Motion planning is tricky – point the vehicle in the wrong direction, and it will dive and crash. We'll introduce state-of-the-art motion planning methods to deal with this.
- We'll use visual odometry (i.e., using computer vision to measure travelled distances) to determine the drone's position and orientation.
- Planning includes dealing with the dynamics of the drone's motion. We'll deal with this using trajectory optimization methods.

Class Schedule Revisited

Projects will focus on specific aspects of the problems associated to one of the modules.

Each project will be assigned on a Tuesday, and will be due at midnight two weeks later.

Week	Month	Date	Day	Module	Topic	Slides	Reading	Quizzes	Projects
1	Aug	23	Tue		Introduction	-	-		Intro to Python
1	Aug	25	Thu	Sorting	State, probability, and actions	-	-		
2	Aug	30	Tue	Sorting	Sensor Models	-	-		P1: Probability and Decision Theory
2	Sep	1	Thu	Sorting	Bayes Law	-	-		
3	Sep	6	Tue	Sorting	Decision Theory	-	-		
3	Sep	8	Thu	Vacuum	Probabilistic Actions	-	-	Quiz 1	
4	Sep	13	Tue	Vacuum	Inference in HMMs	-	-		P2: HMM and MDP
4	Sep	15	Thu	Vacuum	Inference, MDPs	-	-		
5	Sep	20	Tue	Vacuum	MDPs	-	-		
5	Sep	22	Thu	Logistics	Omnidirectional wheels	-	-	Quiz 2	
6	Sep	27	Tue	Logistics	LIDAR sensing	-	-		P3: Particle Filters
6	Sep	29	Thu	Logistics	Monte Carlo Localization	-	-		
7	Oct	4	Tue	Logistics	Planning	-	-		
7	Oct	6	Thu	Duckiebot	Differential Drive	-	-	Quiz 3	
8	Oct	11	Tue	Duckiebot	Cameras and Image Processing	-	-		P3 due
8	Oct	13	Thu	Duckiebot	Computer Vision 101	-	-		
9	Oct	18	Tue	Fall Break	Fall Break	-	-		
9	Oct	20	Thu	Duckiebot	Inference with Deep Nets	-	-		
10	Oct	25	Tue	Duckiebot	Deep Learning	-	-		P4: Deep Learning
10	Oct	27	Thu	Vehicles	Autonomous Vehicles and SE(2)	-	-	Quiz 4	
11	Nov	1	Tue	Vehicles	Ackerman steering	-	-		
11	Nov	3	Thu	Vehicles	SLAM with 3D LIDAR	-	-		
12	Nov	8	Tue	Vehicles	Planning for Avs	-	-		P5: Pose SLAM and LIDAR
12	Nov	10	Thu	Drone	SE(3)	-	-	Quiz 5	
13	Nov	15	Tue	Drone	RRTs for motion planning	-	-		
13	Nov	17	Thu	Drone	Visual Odometry	-	-		
14	Nov	22	Tue	Drone	Trajectory Optimization	-	-		P6: RRT and/or Trajectory Optimization
14	Nov	24	Thu	Thanksgiving	Thanksgiving	-	-		
15	Nov	29	Tue	Extra	Guest Lecture TBD	-	-	Quiz 6	
15	Dec	1	Thu	Extra	Guest Lecture TBD	-	-		
16	Dec	6	Tue	Extra	Guest Lecture TBD	-	-		P6 due, all quizzes graded!

Class Schedule Revisited

Projects will focus on specific aspects of the problems associated to one of the modules.

There will be a quiz at the end of each module.

The module ends on a Tuesday, and the quiz will be in class the following Thursday.

Week	Month	Date	Day	Module	Topic	Slides	Reading	Quizzes	Projects
1	Aug	23	Tue		Introduction	-	-		Intro to Python
1	Aug	25	Thu	Sorting	State, probability, and actions	-	-		
2	Aug	30	Tue	Sorting	Sensor Models	-	-		P1: Probability and Decision Theory
2	Sep	1	Thu	Sorting	Bayes Law	-	-		
3	Sep	6	Tue	Sorting	Decision Theory	-	-		
3	Sep	8	Thu	Vacuum	Probabilistic Actions	-	-	Quiz 1	
4	Sep	13	Tue	Vacuum	Inference in HMMs	-	-		P2: HMM and MDP
4	Sep	15	Thu	Vacuum	Inference, MDPs	-	-		
5	Sep	20	Tue	Vacuum	MDPs	-	-		
5	Sep	22	Thu	Logistics	Omnidirectional wheels	-	-	Quiz 2	
6	Sep	27	Tue	Logistics	LIDAR sensing	-	-		P3: Particle Filters
6	Sep	29	Thu	Logistics	Monte Carlo Localization	-	-		
7	Oct	4	Tue	Logistics	Planning	-	-		
7	Oct	6	Thu	Duckiebot	Differential Drive	-	-	Quiz 3	
8	Oct	11	Tue	Duckiebot	Cameras and Image Processing	-	-		P3 due
8	Oct	13	Thu	Duckiebot	Computer Vision 101	-	-		
9	Oct	18	Tue	Fall Break	Fall Break	-	-		
9	Oct	20	Thu	Duckiebot	Inference with Deep Nets	-	-		
10	Oct	25	Tue	Duckiebot	Deep Learning	-	-		P4: Deep Learning
10	Oct	27	Thu	Vehicles	Autonomous Vehicles and SE(2)	-	-	Quiz 4	
11	Nov	1	Tue	Vehicles	Ackerman steering	-	-		
11	Nov	3	Thu	Vehicles	SLAM with 3D LIDAR	-	-		
12	Nov	8	Tue	Vehicles	Planning for Avs	-	-		P5: Pose SLAM and LIDAR
12	Nov	10	Thu	Drone	SE(3)	-	-	Quiz 5	
13	Nov	15	Tue	Drone	RRTs for motion planning	-	-		
13	Nov	17	Thu	Drone	Visual Odometry	-	-		
14	Nov	22	Tue	Drone	Trajectory Optimization	-	-		P6: RRT and/or Trajectory Optimization
14	Nov	24	Thu	Thanksgiving	Thanksgiving	-	-		
15	Nov	29	Tue	Extra	Guest Lecture TBD	-	-	Quiz 6	
15	Dec	1	Thu	Extra	Guest Lecture TBD	-	-		
16	Dec	6	Tue	Extra	Guest Lecture TBD	-	-		P6 due, all quizzes graded!

Syllabus
Let’s go to the website

Book

This semester, we'll use a new online "text book" for the course:

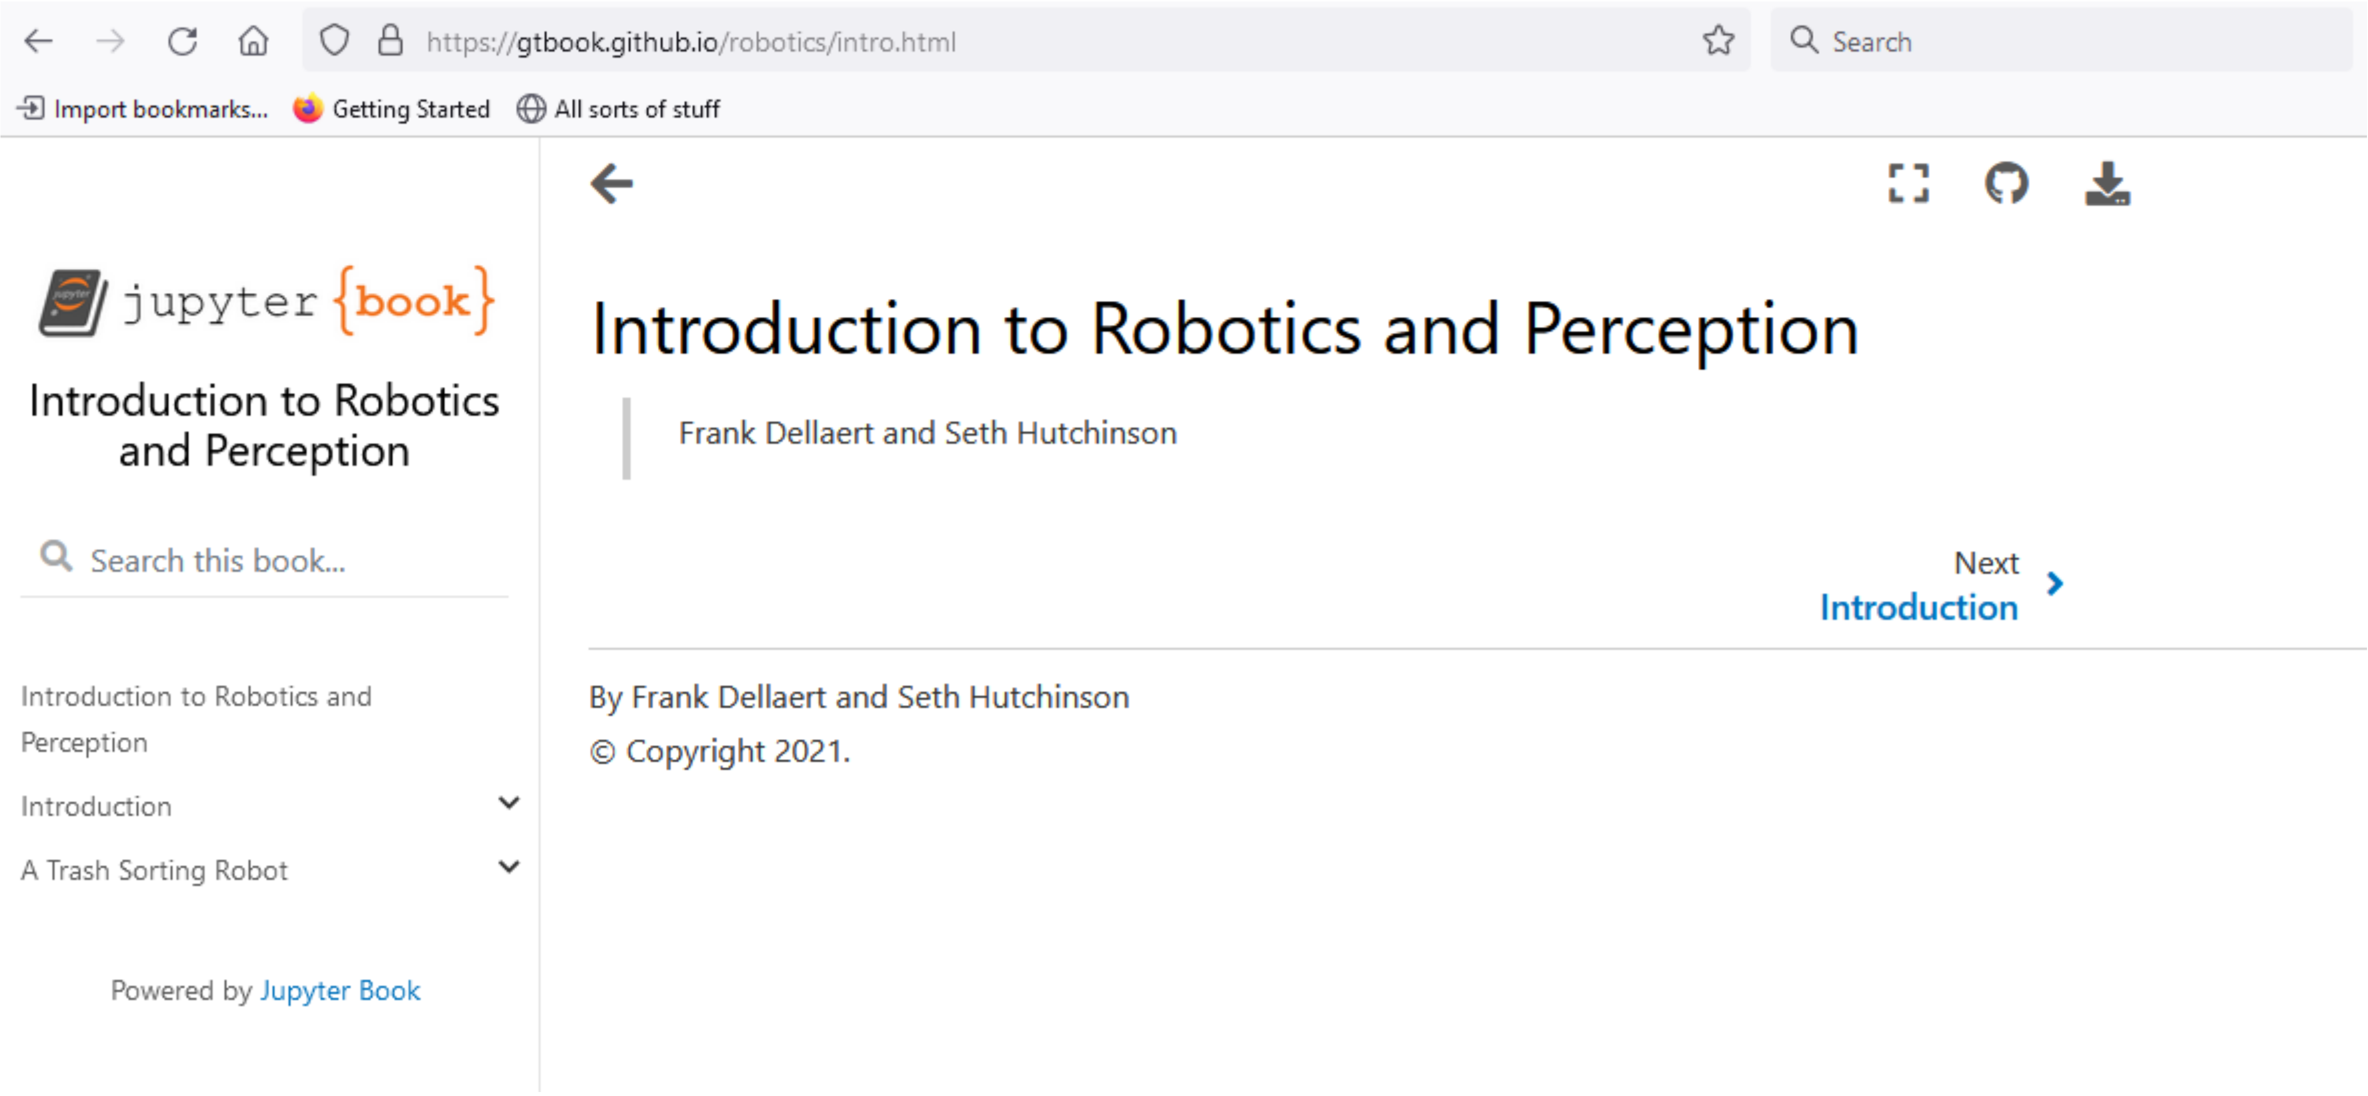

Typical page view for page of text.

jupyter {book}

Introduction to Robotics and Perception

Search this book...

Powered by Jupyter Book

A Trash Sorting Robot

> :A simple robot that sorts trash into appropriate bins can be used to illustrate most all of the concepts that we will teach in this class.

In this chapter we introduce the main concepts that will be covered in this book using a simple trash sorting robot as an example. The robot is fed items of trash, one at a time, by a conveyor belt. Depending on the *category of the item* (state), e.g., paper, metal, etc., the robot's task is to *move the item* (actions) into the appropriate bin. Sensors are used to *measure certain characteristics* (sensing) of a given item of trash, and these sensor measurements are then used to *infer the item's category* (simple perception). Using these inferences, the robot *chooses an action* (planning) to place the item in an appropriate bin. The robot can *refine its model of the world based on data* acquired during operation (learning), and this new model can be used to improve operation over time.

For now, we will abstract away many details of real-world robotics problems. We will not consider the geometric or physical aspects of robot motion, and we will assume that the robot is able to move items from the conveyor belt to the commanded bin without error. On the other hand, we will explicitly consider uncertainty associated to the problem of classifying items of trash into their appropriate categories, and how this uncertainty affects decision making when choosing actions (i.e., destination bins for items of trash). This will allow us to introduce some basic concepts from probability and decision theory that will be used throughout the book.

Previous
6. Learning

Next
1. Modeling the World State

By Frank Dellaert and Seth Hutchinson

A typical page might include text and equations.

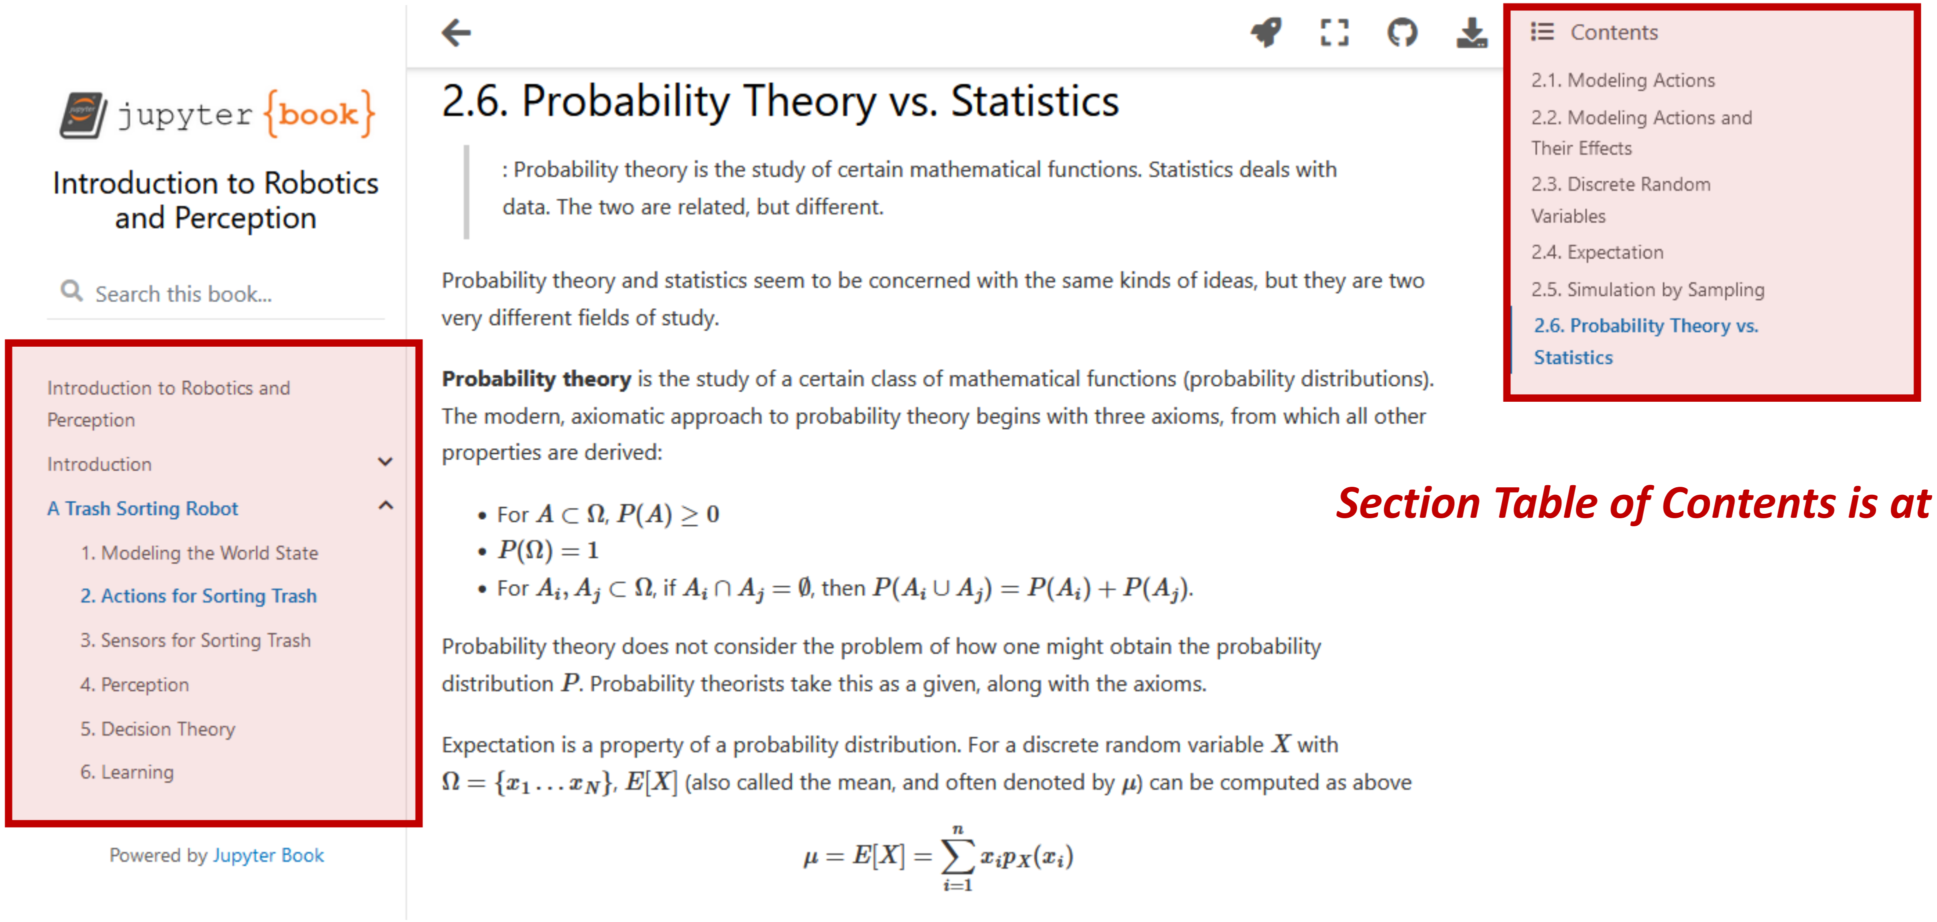

jupyter {book}

Introduction to Robotics and Perception

Search this book...

Powered by Jupyter Book

2.6. Probability Theory vs. Statistics

> : Probability theory is the study of certain mathematical functions. Statistics deals with data. The two are related, but different.

Probability theory and statistics seem to be concerned with the same kinds of ideas, but they are two very different fields of study.

Probability theory is the study of a certain class of mathematical functions (probability distributions). The modern, axiomatic approach to probability theory begins with three axioms, from which all other properties are derived:

- For $A \subset \Omega$, $P(A) \geq 0$
- $P(\Omega) = 1$
- For $A_i, A_j \subset \Omega$, if $A_i \cap A_j = \emptyset$, then $P(A_i \cup A_j) = P(A_i) + P(A_j)$.

Probability theory does not consider the problem of how one might obtain the probability distribution P. Probability theorists take this as a given, along with the axioms.

Expectation is a property of a probability distribution. For a discrete random variable X with $\Omega = \{x_1 \ldots x_N\}$, $E[X]$ (also called the mean, and often denoted by μ) can be computed as above

$$\mu = E[X] = \sum_{i=1}^{n} x_i p_X(x_i)$$

Contents

Section Table of Contents is at the right

Book Table of Contents is at the left

The book is a collection of Jupyter Notebooks.

jupyter {book}

Introduction to Robotics and Perception

Search this book...

Powered by Jupyter Book

Open the Notebook in Colab

Open in Colab

```
%pip install -q -U gtbook
```

```
Note: you may need to restart the kernel to use updated packages.
```

```
from gtbook.discrete import Variables
from gtbook.display import pretty
import numpy as np
import pandas as pd
import gtsam
import plotly.express as px
import plotly.io as pio
pio.renderers.default = "png"
```

2. Actions for Sorting Trash

2.1. Modeling Actions

> : Robots change the world through their actions. Action models capture the salient aspects of those changes.

Robots decide how to act in the world by reasoning about how their actions can be used to achieve their goals, given the current state of the world. At a high level, actions can be represented by symbolic descriptions of their effects (changes that will occur in the world state when the action is executed) and by their preconditions (things that must be true in the current state in order to execute

Contents

The book is a collection of Jupyter Notebooks

jupyter {book}

Introduction to Robotics and Perception

Search this book...

Powered by Jupyter Book

Run code to install libraries

Open in Colab

```
%pip install -q -U gtbook
```

```
Note: you may need to restart the kernel to use updated packages.
```

```
from gtbook.discrete import Variables
from gtbook.display import pretty
import numpy as np
import pandas as pd
import gtsam
import plotly.express as px
import plotly.io as pio
pio.renderers.default = "png"
```

2. Actions for Sorting Trash

2.1. Modeling Actions

> : Robots change the world through their actions. Action models capture the salient aspects of those changes.

Robots decide how to act in the world by reasoning about how their actions can be used to achieve their goals, given the current state of the world. At a high level, actions can be represented by symbolic descriptions of their effects (changes that will occur in the world state when the action is executed) and by their preconditions (things that must be true in the current state in order to execute

Contents

- ***Edit and run python code inline to illustrate concepts.***
- ***By hacking the code, you can try out new ideas, and improve your understanding.***

jupyter {book}

Introduction to Robotics and Perception

Search this book...

Powered by Jupyter Book

2.5. Simulation by Sampling

> :It is easy to demonstrate the relationship between expectation and the average over many trials - simply sample and average!

The code below computes the average cost over N samples for a specified action. Try various values for N, and notice that as N increases, the average tends to be an increasingly better approximation of the expected cost.

```
# Sample N times, and evaluate the cost of executing the given action:
total_cost = 0
N = 100
action = 0
for i in range(N):
    category = category_prior.sample()
    total_cost += cost[action, category]
print(total_cost/N)
```

```
3.36
```

For example, one experiment with 100 samples yielded:

```
cost_estimate = [3.14, 0.6, 4.01, 1.0]
```

Questions?