SPECIAL TOPICS COURSE LISTINGS FOR ACADEMIC YEAR 2025-2026

Optimal Control — 2025 Fall A Term

This course introduces students to the fundamentals of optimal control with a focus on applications in robotics. We begin with foundational concepts in convex optimization, including quadratic programming, which form the basis for many control algorithms. Building on this, students will learn classical optimal control methods such as the Linear Quadratic Regulator (LQR) and Model Predictive Control (MPC). The course then extends to nonlinear MPC, with applications to trajectory planning and control in robotic systems. Finally, we explore emerging learning-based approaches that combine data-driven modeling with optimal control principles, highlighting their potential and limitations in robotics contexts. Throughout the course, theoretical understanding is reinforced with practical examples and assignments involving real or simulated robotic platforms.

Prerequisite: Consent of Instructor

Deep Learning for Perception — 2025 Fall B Term

This course exposes the students to the mathematical foundations of deep learning applied to images. Perception stacks in state-of-the-art robots are rapidly adapting the latest advancements in deep learning due to their efficacy and high accuracy. These deep learning-based methods are also accelerable using parallelized hardware such as GPUs that can enable low latency operations of complex tasks such as real-time scene segmentation. The students will be trained in formulation, development, and implementation of deep learning solutions for common computer vision problems in the context of robot perception. The course will cover advanced and state-of-the-art topics such as sim2real, adversarial attacks on neural networks, vision transformers, and diffusion models. Additional topics explored in this course include image formation, linear classifiers, neural networks and backpropagation, Convolutional Neural Networks (CNNs), CNN architectures, data generation for sim2real, black- and white-box attacks on neural networks as applied to build state-of-the-art robotic stack. Students will gain knowledge about the considerations required to enable a robotic system with the state-of-the-art deep learning toolkit. The course is designed to balance theory with applications through projects.

Recommended Background: Proficiency in programming, preferably Python, MA 1024, MA 2071/20772, MA 2621/2631

Social AI and Robotics — 2025 Fall B Term

This course introduces the rapidly evolving field of human-centered AI and socially assistive robotics. The course is structured in two complementary modules that build upon each other to provide a comprehensive understanding of how AI and robotics can be designed to support human social, emotional, and other healthcare needs. The first half of the course focuses on human-centered AI, covering foundational technologies including large language models (LLMs), affective computing, emotional AI, social AI, and AI personalization techniques that adapt to individual user needs and preferences. The second half covers socially assistive Robotics, examining how these AI and multimodal interaction technologies are implemented in robotic systems designed to provide support, companionship, and assistance to humans. *(continued)*

Topics include social human-robot interaction, design principles for socially assistive robots, research methodologies specific to this interdisciplinary field, and current frontiers in social robotics research.

Prerequisite: Consent of Instructor

Advanced Robot Navigation - Online 2025 Fall Semester

The Advanced Robot Navigation course delves into the sophisticated techniques and algorithms used by autonomous robots to navigate complex, real-world environments with precision and efficiency. Building upon fundamental concepts in robotics and navigation, this course equips students with advanced knowledge and practical skills necessary to develop cutting-edge navigation systems for autonomous robots operating in diverse scenarios. The curriculum begins with a review of foundational principles in robot path and motion planning including RRT*, techniques such as probabilistic methods, Simultaneous Localization and Mapping (SLAM), and multi-sensor fusion. Students will deepen their understanding of these concepts through hands-on exercises and simulations.

Prerequisite Courses: RBE 500 or equivalent.
Recommended Background: RBE 549

Hands-On Autonomous Aerial Robotics 2025 Fall Semester

This is an advanced course on autonomy stack for aerial robots. A unique aspect of this course is its hands-on aspect. Our goal is to train the students to develop real-time perception, planning, control, and AI algorithms for the realization of autonomous aerial systems. The course follows a series of projects that will finally build up to an autonomous drone race on an obstacle course in which the students will compete. This course will explore rigid body transformations, attitude estimation, Bayesian filter, Linear and Unscented Kalman filters, Path and Motion planning, Deep learning, sim2real, Camera models, Gaussian Mixture Models, Image Processing, Visual feature detection and tracking, Projective Geometry, Optical Flow, Stereopsis, Quadrotor Dynamics and Controls, Structure from Motion/SLAM.

Prerequisite: Consent of Instructor

Fundamentals of Artificial Intelligence and Robotics for Autonomous Vehicle Applications – Online 2025 Fall Semester

Fundamentals of **A**rtificial **I**ntelligence and **R**obotics for **A**utonomous **V**ehicle Applications, short for **FAIR-AV**, is a course to introduce you to the interaction of AI, robotics and autonomous vehicles (AVs). The goal of the course is to introduce you to a variety of systems used to implement AVs. These systems and their associated technologies greatly impact how well-automated driving functions can perform in real-world scenarios. More specifically, they are enabling an AV to detect and predict obstacles/objects, plan its motion, and make intelligent decisions to achieve safe and pleasant driving for occupants inside the vehicle. While FAIR-AV will discuss the applicability and challenges of AI and robotics used in wheeled vehicles, we will start with the motivation to bring human-like intelligence to ground vehicles which normally involves the shared control between human drivers and electronically controlled systems. The digital driver in the autonomous vehicle aims to replicate human intelligence, hence the need for artificial intelligence (AI). *(continued)*

After those introductory discussions, the course will introduce to you the basics of automotive system functions and control systems, the necessary mechatronics components, the software modules, and the system engineering aspects. All those teachings will only touch the rudiment knowledge without deeply diving into individual topics on how those tasks are implemented, which will be taught in follow-up courses.
FAIR-AV will use open-sourced software packages and tools as supplementary material if you would like to dive into how various tasks are realized at the algorithm, coding, and software level. The lectures, handouts, homework assignments, and projects are designed to prepare you to grasp the essence of AV to help you to find your passion on individual topics you would like to pursue or to become a system engineer in the AV industrialization.

Prerequisite: RBE 500

Vision-based Robotic Manipulation 2025 Fall Semester

This course focuses on the role of visual sensing in robotic manipulation. It covers fundamental manipulation concepts such as mathematical grasp formulations, taxonomies, and grasp stability metrics. Various grasp planning strategies in the literature are studied. 2D and 3D vision-based control algorithms are covered. Point cloud processing techniques that allow object detection, segmentation, and feature extraction are studied and implemented. Students will integrate all of these aspects to design the whole vision-based robotic manipulation pipeline.

Recommended background: Knowledge of robot kinematics, wrench spaces, and rigid body transformations as presented in RBE 3001 or RBE 500. Familiarity with robotic simulation software as presented in RBE 3002 or RBE 500.

The course is cross-listed with RBE 4540 in the A-term. Graduate students will apply material from the course to projects in the B-term.

Reinforcement Learning - Online 2026 Spring Semester

This course will provide a solid introduction to the field of Reinforcement Learning (RL). Students will learn about the core challenges and approaches including Markov decision processes, model based, model free RL, on-policy and off-policy learning, and approximate solution techniques. Through a combination of lectures and coding assignments, students will become well versed in key ideas and techniques in RL and its application in robotic systems. To get students familiarized with the state-of-the-art RL algorithms in robotics, research papers are provided, and students are required to give a presentation about the papers. In addition, an end of the term team project would allow the students to apply mastery of the subject to a real-world robotics application.

Prerequisites: A probability course is required, as well as proficiency in Python. RBE 500/Foundations of Robotics and basic knowledge of neural networks preferred but not required.