AE 6512 Syllabus

Mathematical principles of planning and decision-making for autonomy- 3 Credits

General Information

Description

The course will introduce the students to the mathematical tools and the theory for formulating and solving problems that involve high-level decision-making for controlling and planning of autonomous systems with an emphasis on aerospace applications.

Pre- &/or Co-Requisites

Corequisites: AE6530

Course Goals and Learning Outcomes

The course will introduce the students to the mathematical tools and the theory for formulating and solving problems that involve high-level decision-making for controlling and planning of autonomous systems emphasizing problems of interest to aerospace engineering. Upon completion of this course the students will be able to:

- Understand Model-based vs model-free decision-making
- Develop graph search algorithms to solve a variety of path-planning problems
- Use Markov decision processes (MDP) to model stochastic sequential decision-making problems
- Use Dynamic programing (DP) and reinforcement learning to solve sequential decision-making problems
- Understand elements of game theory for decision-making in the presence of multiple intelligent agents

Course Requirements & Grading

Note: Graded components of a course may vary with each offering. The example below is typical but subject to change.

Description of Graded Components

Final Grade = Class Participation 10% + Homework 40% + Final Exam/Project 50%

Grading Scale

Your final grade will be assigned as a letter grade according to the following scale:

A	85-100%
B	70-84%
C	51-69%
D	46-50%
F	0-45%

Topics Covered

Note: The exact topics covered in a course may vary with each offering. The example below is typical but subject to change.

- Brief overview of optimal control
- Hamilton-Jacobi theory, the principle of optimality
- LQR, DDP/iLQR, MPC control design methods
- Reactive vs. deliberative agents
- Model-based vs model-free decision-making
- Search space representations, state partitioning, cell and Voronoi decompositions
- Elements of graph theory
- Uninformed graph search algorithms, depth-first and breadth-first search
- Informed graph search algorithms, A*, Dijkstra, D*, LPA*
- Sampling-based methods: PRM, RRT, RRT*
- Probabilistic inference, graphical models, Hidden Markov Models (HMMs)
- Decision-making under uncertainty, stochastic search
- Markov decision processes (MDP), partially observable MDP (POMDP)
- Dynamic programing, value and policy iteration
- Approximate dynamic programing
- Reinforcement learning and inverse reinforcement learning
- Adversarial agents, stochastic games, competitive MDPs
- Nash and correlated equilibria
- Games of incomplete information, Bayesian games, bounded rationality
- Differential games, games of pursuit and evasion

Course Materials

Note: Course materials may vary with each offering. The example below is typical but subject to change.

Textbook

There is no required textbook. A clear, detailed set of lecture notes should be sufficient to follow the material in the class.

Course notes

A concise set of notes will be shared via Canvas in the form of slides