E-MDS130 Reinforcement Learning

Introduction

Artificial Intelligence in this era has become synonymous with Supervised and Unsupervised Learning. Supervised learning is best suited for cases that have a vast set of examples of inputs and desired outputs and the objective is to learn based on such examples in order to generate output from some future, currently unseen input. Text classification, Image Classification, Object location, Regression problems, and Sentiment analysis are areas where supervised learning is extensively used. Whereas, Unsupervised learning aims to discover some hidden structure of the data without the need to have the specific distinction in the input and output values. Such learning techniques are commonly used for clustering of data that tries to combine data items into a set of clusters revealing relationships in data.

Reinforcement learning lies somewhere in between supervised and unsupervised forms of learning techniques. On one hand, it builds on established methods of supervised learning for function approximation, stochastic gradient descent, and backpropagation to learn data representation, however, on the other hand, it does not require supervision in order to discover hidden patterns and relationships in data. Reinforcement learning primarily focuses on the problem of automatic learning of optimal decisions over time in a complex environment by building on advances in computer science, behavioral psychology, and neuroscience. Due to its flexibility and generality, the field of RL is developing very quickly and attracting lots of attention, both from researchers who are trying to improve existing methods or create new methods and from practitioners interested in solving their problems in the most efficient way.

The target group for this course would be professionals and students working or interested in areas of artificial intelligence, machine learning, game theory, control theory, operations research, information theory, simulation-based optimization, multi-agent systems, swarm intelligence, and statistics.

Learning Outcomes

This course would introduce AI and Optimization in a fun, easy, interesting, immersive, and hands-on way. Optimization problems are becoming essential across multiple disciplines. The skills gained through this course would allow the use of efficient optimization strategies in work processes. These could include optimization of complex machine learning models making them more efficient, creating exploratory models that without training can evaluate a situation and gradually make positive decisions, explore financial data to discover patterns that lead to beneficial outcomes, and others.

Content

Concepts covered in this course would provide relevant theoretical and hands-on programming knowledge. Every topic is demonstrated using easy-to-understand real-world examples. The following topics would be covered during the course duration:

- Topic 1: Reinforcement Learning - an introduction
- Topic 2: Course Materials, Supplementary Resources, and Development Environment

- Topic 3: Tabular Methods
    - Multi-Armed Bandit
    - Markov Decision Processes
    - Cross-Entropy Method
- Topic 4: Dynamic Programming
- Topic 5: Monte-Carlo & Temporal Difference and Q-Learning
- Topic 6: Policy Gradients
- Topic 7: The Actor-Critic Method
- Topic 8: Deep Q-Network – an Overview
- Topic 9: Further Exploration

Form of Work

- A total of 10 – 12 lectures will cover the course syllabus
- Every lecture would be covered through bite-size recorded videos.
- Three hands-on ***use-case based*** assignments will be presented to the students
- Four live sessions would be conducted to discuss the solutions for each assignment

Prerequisites

- Good programming knowledge
- Knowledge of basic algebra, probability, and statistics
- Python Programming Knowledge
- Understanding of Numpy, Matplotlib

Exam / Assessment

- Individual home exam – All aids
- 75 Multiple Choice Questions
- Exam Date: xx.xx.2022
- Grade: A- F

Conditions for taking the Exam / Assessment

- The three assignments are mandatory – approved / not approved
- Submissions should provide coding solutions to the respective problems with proper documentation.

Topic Evaluation

The course will follow student evaluation procedures established by UiS EVU.

Literature

- Sutton, Richard S., and Andrew G. Barto. Reinforcement learning: An introduction. MIT Press, 2018.
- Lapan, Maxim. Deep reinforcement learning hands-on. Packt Publishing, 2020.