ITMD 524 SYLLABUS

ITMD 524 Applied Artificial Intelligence and Deep Learning

Hours: 3 credit hours / 45 contact hours

Instructor: Yong Zheng

Textbook, title, author, and year:
(all are optional but are strongly recommended)

Artificial Intelligence: A Modern Approach (4th Edition). Stuart Russell, Peter Norvig. 2020

Machine Learning: The Art and Science of Algorithms that Make Sense of Data, Peter Flach. 2012

Deep Learning. Ian Goodfellow, Yoshua Bengio, Aaron Courville. 2015

Specific course information

a. **Catalog description:** Artificial Intelligence (AI) is being used extensively to solve real-world complex problems. This course will deliver concepts and skills in both classical AI and modern AI. The classical AI refers to the fundamental knowledge in AI, such as search, logic, planning, uncertainty, game theory, Markov models, etc. Modern AI, by contrast, will be concentrated on machine learning and deep learning techniques, especially their applications in NLP, object recognition, recommender systems, etc. Students will learn how to use Python to solve specific AI problems.
b. **Prerequisites:** ITMD 514 or ITMS 514

Specific goals for the course

a. **Program Educational Outcome:** Deliver optimal technical and policy technology solutions for the problems of business, industry, government, non-profit organizations, and individuals in each student's particular area of focus.
b. **Course Outcomes:** At the completion of the course, each student will have the capability to deal with real-world AI problems, such as search, logics, uncertainty, optimization, etc. Particularly, students should be familiar with the knowledge and concepts in machine learning and deep learning, and they are able to utilize the libraries in Python to deal with deep learning applications, such as NLP, object recognition, recommender systems, etc.
c. **Course Student Outcomes:**
Upon successful completion of the course the student should be able to do the following:
   - Describe what AI is and why it is so useful
     - Describe real-world AI problems and applications
   - Recall and distinguish between related terms: artificial intelligence, data mining, machine learning, etc
   - Explain and employ traditional AI problems and techniques, including search, logics, uncertainty, probability, markov model, game theory, etc
   - Perform popular data preprocessing tasks: data selection, data cleaning, data transformation, etc
   - Describe the difference between different machine learning categories, including supervised learning, unsupervised learning, semi-supervised learning, solutions
   - Perform and evaluate different AI solutions using Python
   - Use the deep learning framework TensorFlow
     - Use TensorFlow to solve AI problems and build AI applications

Topics to be covered

a. Intro: Data Mining, Machine Learning & AI
b. Traditional AI: Search 1
c. Traditional AI: Search 2
d. Traditional AI: First-Order Logics
e. Traditional AI: Bayes Nets and Uncertainty
f. Traditional AI: Uncertainty and Probability
g. Traditional AI: Markov Model & Game Theory
h. Midterm Exam and Intro to Machine Learning
i. Modern AI: Machine Learning & Optimization
j. Modern AI: Neural Networks & Deep Learning
k. Modern AI: ANN, CNN, RNN
l. Modern AI: ANN, CNN, RNN with TensorFlow
m. Applications: NLP
n. Applications: Object Recognition
o. Applications: Recommender Systems
p. Exams or Project Presentations