Introduction to AI – Probabilistic Reasoning and Decision Making

Department: Fudan International Summer Session

<table>
<tr><td>Course Code</td><td colspan="3">ECON170026</td></tr>
<tr><td>Course Title</td><td colspan="3">Introduction to AI – Probabilistic Reasoning and Decision Making</td></tr>
<tr><td>Credit</td><td>2</td><td>Credit Hours</td><td>36+3 tutorial hours
(one credit hour is 45 minutes)</td></tr>
<tr><td>Course Nature</td><td colspan="3">□Specific General Education Courses □Core Courses ☑General Education Elective Courses □Basic Courses in General Discipline □Professional Compulsory Courses □Professional Elective Courses □Others</td></tr>
<tr><td>Course Objectives</td><td colspan="3">• Describe and use different probabilistic models including Bayes Nets and EM algorithm
• Apply probabilistic models to solve real-world problems
• Design specific models for AI tasks
• Perform inference using probabilistic models
• Prove relationships between probabilities under different models
• Implement core algorithms of different models
• Describe how agents learn from data using maximum likelihood learning
• Identify ethical concerns related to AI</td></tr>
<tr><td>Course Description</td><td colspan="3">This course will introduce students to the probabilistic and statistical models at the heart of modern artificial intelligence. Possible topics to be covered include: probabilistic methods for reasoning and decision-making under uncertainty; inference and learning in Bayesian networks</td></tr>
<tr><td colspan="4">Course Requirements:
Prerequisites are elementary probability, statistics, linear algebra, and calculus, as well as basic programming ability in some high-level language such as C, Java, Matlab, R, or Python. Programming assignments are completed in the language of the student's choice. Students should satisfy the pre-requisites before enrolling this class as AI classes are math heavy.</td></tr>
<tr><td colspan="4">Teaching Methods:
Lectures</td></tr>
</table>

Instructor's Academic Background:

Dr. Cao Yingjun received his Ph.D. in Computer Engineering from Duke University (Durham, NC, USA). His primary research interest is Computer Science Education with a focus on collaborative learning. He is also involved in more traditional research on network data analysis and distributed learning. Dr. Cao has been teaching in the Department of Computer Science & Engineering at the University of California, San Diego since 2015.

Email: yic242@eng.ucsd.edu

Course Schedule (Tentative, subject to change)

Day 1: Course overview, Modeling uncertainty, review of probability.
Day 2: Examples of probabilistic reasoning.
Day 3: Belief networks: from probabilities to graphs.
Day 4: Conditional independence, d-separation.
Day 5: Inference in polytrees and loopy networks.
Day 6: Learning, maximum likelihood estimation.
Day 7: Naive Bayes and Markov models.
Day 8: Latent variable models, EM algorithm.
Day 9: Examples of EM algorithm.
Day 10: Estimating Noisy-OR parameters with EM
Day 11: Other AI topics (such as HMM or Q-learning)
Day 12: Final Exam

The design of class discussion or exercise, practice, experience and so on:

Exercise and practice

Grading & Evaluation:

Your final grade will be determined via the following percentages:
Lecture participation points: 10%
Homework: 45%
Final: 45%

Important grading policies:

- You must score at least 55% on the final exam to pass the course. If you score lower than 55% on the final, you will receive an F for the course, regardless of your overall average.
- You must score at least 55% overall for the homework assignments.
- All homework should be done individually.
- According to Fudan University's policy, there is a threshold on the percentage of students who may receive A or A- in a class. Please keep this policy in mind.
- There will be no make-up exam

Teaching Materials & References (Including Author, Title, Publisher and Publishing time):

- Artificial Intelligence: Foundations of Computational Agents, 2nd ed. by Poole and Mackworth. An online version of this textbook can be found on the publisher's website: https://artint.info/2e/html/ArtInt2e.html
- Artificial Intelligence: A Modern Approach, 3rd ed by Russell and Norvig.