האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

The Hebrew University of Jerusalem

Syllabus

PROBABILISTIC METHODS IN ARTIFICAL INTELLIGENCE - 67800

Last update 11-09-2017

HU Credits: *4*

Degree/Cycle: *1st degree (Bachelor)*

Responsible Department: *computer sciences*

Academic year: *0*

Semester: *2nd Semester*

Teaching Languages: *Hebrew*

Campus: *E. Safra*

Course/Module Coordinator: *Gal Elidan*

Coordinator Email: *galel@cs.huji.ac.il*

Coordinator Office Hours: *By appointment*

Teaching Staff:
Prof
Mr. Eitan Richardson

Course/Module description:
The course will cover the framework of probabilistic graphical models, a core machine learning paradigm for making predictions in uncertain environments. In the last part of the course (~4 weeks) we will transition from prediction to action models and present the fundamentals of Reinforcement Learning.

Course/Module aims:
To develop the probabilistic graphical model representation, to present algorithms for inference and learning using these models, to demonstrate the applicability of these model for real problems; Present the fundamentals of Reinforcement Learning with and without explicit modeling of the underlying probabilistic model.

Learning outcomes - On successful completion of this module, students should be able to:
Prove basic claim in the field of probabilistic graphical models and reinforcement learning, adapt tools to new scenarios, demonstrate understanding of learned algorithms

Attendance requirements(%):
0

Teaching arrangement and method of instruction: Lecture, TA class

Course/Module Content:
The probabilistic graphical representation: Bayesian and Markov networks, algorithms for example inference, approximate inference using propagation based, variational and sampling algorithms, parameter and structure learning. Reinforcement learning: Markov decision processes, value iteration, policy iteration, q-learning.

Required Reading:
None

Additional Reading Material:

Probabilistic Graphical Models: Principles and Techniques by Daphne Koller and Nir Friedman

Reinforcement Learning: An Introduction by Richard Sutton and Andrew Burton

Course/Module evaluation:

End of year written/oral examination 50 %
Presentation 0 %
Participation in Tutorials 0 %
Project work 0 %
Assignments 50 %
Reports 0 %
Research project 0 %
Quizzes 0 %
Other 0 %

Additional information:

The average of the best n-1 of n programming exercises and m-1 of m theoretical exercises will account for 50% of the grade, while the final exam will account for the other 50%.
To pass the course, a passing grade is needed both in the exam and in the exercises.

The course is primarily aimed at graduate students and it is recommend not to take it before the third year if taken as part of an undergraduate degree.