האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

The Hebrew University of Jerusalem

Syllabus

INTRODUCTION TO MULTIAGENT SYSTEMS - 67715

Last update 27-04-2015

HU Credits: *2*

Degree/Cycle: *2nd degree (Master)*

Responsible Department: *computer sciences*

Academic year: *3*

Semester: *2nd Semester*

Teaching Languages: *English*

Campus: *E. Safra*

Course/Module Coordinator: *Prof Jeff Rosenschein*

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

Coordinator Office Hours: *Tuesdays, 10:30am-11:00am*

Teaching Staff:

Prof Jeff Rosenschein

Course/Module description:

Multiagent systems have emerged as an important area of research and development in information technology. A multiagent system is one composed of multiple interacting software components known as agents, which are typically capable of cooperating to solve problems that are beyond the abilities of any individual member. This course explores, through lectures and readings, theoretical agent models, architectures (reactive and BDI), multiagent interactions, interagent communication languages, agent design methodologies, and agent applications.

Here is the syllabus, in detail:

1. Lecture 1 – Introduction to Multiagent Systems
a. Trends in computing
i. ubiquity
ii. interconnection
iii. intelligence
iv. delegation
v. human-orientation
b. Definitions of "agent" and Multiagent Systems
c. Example systems
d. Some views of the field

2. Lecture 2 – Intelligent Agents
a. What is an agent?
i. Reactivity
ii. Proactiveness
iii. Balancing Reactive and Goal-Oriented Behavior
b. Environments
i. Accessible vs. inaccessible
ii. Deterministic vs. non-deterministic
iii. Episodic vs. non-episodic
iv. Static vs. dynamic
v. Discrete vs. continuous
c. Agents as Intentional Systems
d. Abstract Architecture for Agents
i. Purely Reactive Agents
ii. Agents with State
iii. Utility functions over states
iv. Utility functions over runs
v. Bounded optimal agents
vi. Predicate task specifications
vii. Task environments
viii. Achievement and Maintenance Tasks

ix. Agent Synthesis

3. Lecture 3 – Deductive Reasoning Agents
a. Agent Architectures
b. Symbolic Reasoning Agents
c. Deductive Reasoning Agents
d. Planning systems
i. The Blocks World
ii. Green's Method
iii. The Frame Problem
iv. Frame Axioms
e. AGENT0 and PLACA
f. METATEM and Concurrent METATEM

4. Lecture 4 – Practical Reasoning
a. Intentions in Practical Reasoning
b. Planning agents (again)
c. Means-end reasoning
d. STRIPS
e. Deliberation
i. Commitment Strategies
ii. Intention Reconsideration
f. BDI Theory and Practice
g. BDI Logic
i. Axioms of KD45
ii. CTL Temporal Logic
iii. CTL Notation*
h. Implemented BDI Agents
i. IRMA
ii. PRS
iii. HOMER
1. Comparison with SHRDLU

5. Lecture 5 – Reactive and Hybrid Architectures
a. Brooks – behavior languages
i. The Subsumption Architecture
ii. Layered Control
b. Steels' Mars Explorer
c. Situated Automata
i. RULER
ii. GAPPS
d. Advantages of Reactive Agents
e. Limitations of Reactive Agents
f. Hybrid Architectures
i. Horizontal Layering
ii. Vertical Layering (1-pass control)

iii. Vertical Layering (2-pass control)
iv. Ferguson – TOURINGMACHINES
v. Müller –InteRRaP

6. Lecture 6 – Multiagent Interactions
a. Multiagent Systems
b. Utilities and Preferences
c. Multiagent Encounters
d. Rational Action
e. Payoff Matrices
f. Dominant Strategies
g. Nash Equilibrium
h. Competitive and Zero-Sum Interactions
i. The Prisoner's Dilemma
j. Axelrod's Tournament
k. The Game of Chicken and other symmetric 2 x 2 games

7. Lecture 7 – Reaching Agreements – Auctions, Negotiation, Argumentation
a. Mechanisms, Protocols, and Strategies
b. Mechanism Design
c. Auctions
i. Auction Parameters
1. Goods can have:
a. private value
b. public/common value
c. correlated value
2. Winner determination may be:
a. first price
b. second price
3. Bids may be:
a. open cry
b. sealed bid
4. Bidding may be:
a. one shot
b. ascending
c. descending
ii. English Auctions
iii. Dutch Auctions
iv. First-Price Sealed-Bid Auctions
v. Vickrey Auctions
vi. Lies and Collusion
d. Negotiation
i. Task Oriented Domains
1. The Postmen Domain
2. The Database Domain
3. The Fax Domain

ii. State Oriented Domains
1. The Slotted Blocks World
iii. Worth Oriented Domains
1. The Multiagent TileWorld
iv. Task Oriented Domains Defined
1. Deals
2. The Negotiation Set
3. Negotiation Protocols
4. The Monotonic Concession Protocol
5. The Zeuthen Strategy
6. Deception
a. Hiding information
b. Phantom information
c. Decoy information
7. Negotiation over Mixed Deals
8. Sub-Additive TODs
9. Incentive Compatible Mechanisms
10. Concave TODs
11. Modular TODs
e. Argumentation
i. Logic-based Argumentation
ii. Attack and Defeat
iii. Abstract Argumentation

8. Lecture 8 – Coalitions, Voting Power, and Computational Social Choice
a. Forming Coalitions: Coalitional Games
i. Coalition Structure Generation
ii. Solving the optimization problem of each coalition
iii. Dividing the Benefits
b. Formalizing Cooperative Scenarios
i. The Core
ii. Objections
iii. The Shapley Value
c. How to Represent Characteristic Functions
i. Induced Subgraph
ii. Weighted Voting Games
iii. Marginal Contribution Nets
d. Voting Power
i. Shapley-Shubik Index
ii. Weighted Voting Bodies
iii. Power of Voting Blocs
iv. Committees
e. Preference Aggregation and Social Choice Theory
f. Ordinal voting methods
g. Arrow’s Impossibility Theorem
h. Gibbard–Satterthwaite Theorem

i. Manipulations
j. Voting Protocols
i. Sequential Pairwise Voting
ii. Plurality Voting
1. Majority Graphs
iii. Plurality with Runoff
iv. The Borda Count
v. Copeland
vi. Single Transferable Vote (STV)
k. Voting Criteria
i. Pareto Criterion
ii. Condorcet Winner Criterion
iii. Condorcet Paradox
iv. Monotonicity Criterion

9. Lecture 9 – Agent Communication
a. Speech Acts
b. Plan Based Semantics
c. KQML
d. KIF – Knowledge Interchange Format
e. Criticisms of KQML
f. FIPA
i. Inform and Request

10. Lecture 10 – Working Together
a. Benevolent Agents
b. Task Sharing and Result Sharing
c. The Contract Net
i. Recognition
ii. Announcement
iii. Bidding
iv. Awarding & Expediting
v. Issues for Implementing Contract Net
d. Cooperative Distributed Problem Solving (CDPS)
e. Distributed Sensing
f. The Hearsay II architecture
i. Knowledge Sources
ii. The Blackboard
g. The Distributed Hearsay II architecture
h. The Distributed Vehicle Monitoring Testbed (DVMT)
i. Functionally Accurate/Cooperative (FA/C) systems

<u>Course/Module aims:</u>
To introduce students to the Artificial Intelligence research subarea of Multiagent

Systems, with emphasis on definitions of Intelligent Agents; Deductive Reasoning Agents; Practical Reasoning; Reactive and Hybrid Architectures; Multiagent Interactions; Reaching Agreements using Auctions, Negotiation, Argumentation; Coalitions, Voting Power, and Computational Social Choice; Agent Communication; and Cooperative Distributed Problem Solving.

Learning outcomes - On successful completion of this module, students should be able to:

See course aims

Attendance requirements(%):

100

Teaching arrangement and method of instruction: Frontal lecture. Students are also required to carry out weekly readings, and submit weekly questions about that week's readings.

Course/Module Content:

See course description

Required Reading:

An Introduction to Multiagent Systems. Michael Wooldridge (John Wiley and Sons), second edition, 2009.

Additional Reading Material:

1. Multiagent Systems. Yoav Shoham and Kevin Leyton-Brown (Cambridge University Press), 2009.

2. Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents, Stan Franklin and Art Graesser. Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages, Springer-Verlag, 1996.

3. Agent-Oriented Programming, Yoav Shoham, Artificial Intelligence, Volume 60, 1993, pp. 51-92.

4. Plans and Resource-Bounded Practical Reasoning, M. E. Bratman, D. J. Israel, and M. E. Pollack, Computational Intelligence Journal, Vol. 4, No. 4, pp. 349- 355, 1988.

5. The Belief-Desire-Intention Model of Agency, M. Georgeff, B. Pell, M. Pollack, M.

Tambe, and M. Wooldridge, in J. P. Muller, M. Singh, and A. Rao, editors, Intelligent Agents V, Springer-Verlag Lecture Notes in AI, Volume 1365, March 1999.

6. A Robust Layered Control System for a Mobile Robot, Rodney A. Brooks, IEEE Journal of Robotics and Automation, Vol. 2, No. 1, March 1986, pp. 14- 23.

7. A Situated View of Representation and Control, Stanley J. Rosenschein and Leslie Pack Kaelbling, Artificial Intelligence, volume 73, 1995, pp. 149-173.

8. Chapters 1, 2, and 3 (pages 3-69) of The Evolution of Cooperation, Robert Axelrod, Basic Books, New York, 1984.

9. Automated Negotiation: Prospects, Methods and Challenges, N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra and M. Wooldridge, International Journal of Group Decision and Negotiation, Vol. 10, No. 2, 2001, pp. 199-215.

10. Computational-Mechanism Design: A Call to Arms, Rajdeep K. Dash, Nicholas R. Jennings, and David C. Parkes, IEEE Intelligent Systems, November-December 2003, pp. 40-47.

11. Consenting Agents: Designing Conventions for Automated Negotiation, Jeffrey S. Rosenschein and Gilad Zlotkin, AI Magazine, Volume 15, Number 3, Fall 1994, pp. 29-46.

12. "Designing the Perfect Auction," Hal R. Varian, Communications of the ACM, Volume 51, Number 8, August 2008, pp. 9-11.

13. "Making Decisions Based on the Preferences of Multiple Agents," V. Conitzer, Communications of the ACM, Vol. 53, No. 3, March 2010 , pp. 84- 94

14. Elements of a Plan-Based Theory of Speech Acts, P. R. Cohen and C. R. Perrault, Cognitive Science Vol. 3, No. 3, 1979, pp. 177-212; reprinted in: Readings in Natural Language Processing, Grosz, B.J., Sparck-Jones, K., and Webber, B.L., (eds.), Morgan-Kaufman Publishing Co., Los Altos, 1986; Readings in Distributed Artificial Intelligence, Gasser, L., and Huhns, M., (eds.), Morgan-Kaufman Publishing Co., Los Altos, 1988.

15. Distributed Problem Solving and Planning, E. H. Durfee, in Multiagent Systems: A Modern Approach to Distributed Artifical Intelligence, edited by G. Weiss, pp. 121-164, MIT Press, Cambridge, MA, 1999.

Course/Module evaluation:
End of year written/oral examination 100 %
Presentation 0 %

Participation in Tutorials 0 %
Project work 0 %
Assignments 0 %
Reports 0 %
Research project 0 %
Quizzes 0 %
Other 0 %

Additional information:

There is a midterm exam given, that can be used for 20% of the student's grade, if it is higher than the end-of-semester grade. Students are also required to carry out weekly readings, and submit weekly questions about that week's readings.