CS 7180—Fall 2012
Behavioral Modeling and Decision-making in Artificial Intelligence

Instructor
Professor Amy Sliva
Office: 256 West Village H
Office Hours: W 10:00-11:00am, F 1:30-2:30pm, or by appt.
asliva@ccs.neu.edu

Course Description

Prerequisite:	I will assume you have a basic knowledge of search algorithms (i.e., depth first, breadth first, heuristic, etc.), propositional and first-order logic, probability theory, and basic complexity theory
Credits:	4 credits
Class:	W,F 11:45am–1:25pm, 155 Ryder Hall
Website:	http://www.ccs.neu.edu/course/cs7180f12

In many domains—game playing, economics, politics, etc.—it is necessary to make decisions *strategically*, that is, based on the predicted behavior of someone else. For example, when determining your next move in chess, you must consider the move your opponent just made, as well as what he is likely to do in future turns. In many circumstances, however, making these strategic decisions is extremely difficult (even impossible) for human analysts due to the size, complexity, or uncertainty, of the problem domain.

The field of artificial intelligence has developed many techniques to automate this process of behavioral analysis and decison-making with special attention paid to situations of high uncertainty. This course will survey these methods, providing a look at tools used to model behavior and make strategic decisions. The course will be divided into two main units: 1) behavior representation and modeling, and 2) decision-making under uncertainty. We will study both the basic techniques and algorithms related to topics such as state-space search, logic, constraint satisfaction, and probabilistic graphical models, as well as examine new research that utilizes these tools for predictive modeling and decision-making.

Course Administration

Lecture notes and weekly readings will be available online on the Schedule page. Readings will only be accessible with a valid CCS account and password.

All announcements, lectures, readings, and and a class discussion forum will also be available on the CS7180 Piazza at https://piazza.com/northeastern/fall2012/cs7180.

This course will introduce behavioral models and decision-making in AI through lectures, readings, and discussions of relevant research topics. As this course is designed to be a seminar, reading and discussion will comprise a substantial part of your grade. Throughout the semester we will have student presentations of assigned research papers, so you must come to class prepared to discuss the material and even *teach* your fellow students.

Required Textbook

Artificial Intelligence: A Modern Approach, **3rd Edition** by Stuart Russell and Peter Norvig. (Prentice Hall 2010).
BE SURE TO GET THE 3rd EDITION—it is significantly different (and better) than the 2nd edition!!

Readings and Presentations

Readings will be assigned from *Artificial Intelligence: A Modern Approach* (AIMA), as well as a variety of additional books and research publications. These readings are chosen to provide both a foundation for studying behavior and decision-making in AI, as well as an understanding of current research directions and advances. All readings **NOT** from AIMA will be available on the course website.

One of the most important aspects of research or development in any field is presenting and explaining your findings to your peers and supervisors. As such, throughout the semester you will each have a chance to put on your "professor" hat and present one of the assigned readings to the class. There will be 2 or 3 students per research paper, with a total of 25 (groups of 2) or 35 (groups of 3) minutes. You may divide the content however you choose, use slides (or not), and should thoroughly explain the content of the papers to the class. In addition, you must identify some topics for discussion: potential future research directions, insights into where this research may lead, what challenges they face, or what steps might have to be accomplished before this is widely implemented. Do not just present the conclusions or future work paragraph from the paper, but provide your own thoughts for discussion on the potential of this research.

Assignments

Throughout the semester you will have 3 or 4 assignments to give you practice with the course material. You may collaborate with your classmates, but your final solutions must be your own. If you work with other students, please note their names on your assignment when you turn it in.

Exams

The class includes one midterm examination that will test your understanding of the lecture material as well as some of the research papers discussed in class. A **tentative** date for the exam is:

November 2, 2012

The exact date will be confirmed later, and may vary due to the progress of the course and other factors.

Term Project

A significant portion of this course will be a term project, which will essentially be a miniature version of the research projects you will do repeatedly throughout your career.

The projects will procede according to the follow schedule:

Late September:	Your teams should be formed by this time
October 19:	Term project proposals due by midnight (11:59:59pm)
October 24:	Present your proposals in class (10 minutes per group)
December 3:	Term project reports due by midnight
December 5, 7:	Present your reports in class (20 minutes per group)

The projects will be done in teams of 3. Start thinking **NOW** about who you would like to work with. Some things to think about when forming your teams:

1. Do you feel comfortable with your team?
2. Do their interests and abilities complement yours?
3. Do you think you can depend on them?
4. Do you think you can work well together?

Grading

Your course grades will be determined according to the following percentages:

Assignments	15%
Presentations	15%
Midterm	30%
Term Project	40%

Class attendance and participation in research discussions will also be taken into account in determining the course grade. **Late assignments will not be accepted unless prior arrangements have been made with the instructor.**

Academic Honesty

The individual assignments must be each student's own work. Any group projects assigned must be the work of the students in the group. Plagiarism or copying will result in official University disciplinary review. Security is an important aspect of computer science. Students are expected to protect their work from plagiarists.

Course Evaluations

The Teacher Rating and Course Evaluation (TRACE) survey is a required part of every course at Northeastern University. All students are expected to participate in TRACE at the end of the semster to provide feedback about the course and instructor. This will help us develop better courses for you!

Schedule (Tentative)

The following schedule may vary according to the pace of the course.

September 5: ***Introduction to AI and Strategic Decision-making***
AIMA 1.1

Laura A. McNamara, Timothy G. Trucano, and Charles Gieseler. 2011. *Challenges in Computational Social Modeling and Simulation for National Security Decision Making*. Defense Threat Reduction Agency and Sandia National Laboratory.

Thomas H. Davenport and Jeanne G. Harris. 2005. Automated Decision Making Comes of Age, *MIT Sloan Management Review*, 83-89.

September 7: ***Reasoning in AI: Overview of Propositional and First-order Logic***
AIMA Chapter 7, 8, and 9

September 12: ***Guest speaker—Dr. Glenn Pierce, Institute of Security and Public Policy & School of Criminology and Criminal Justice***

September 14: ***Guest speaker—Professor Thomas Wahl: First-order Logic and Applications to Program Reasoning***
AIMA Chapters 8 and 9

Sign up online for research presentations.

September 19: ***Temporal Logic for Reasoning about Events***
Michael Fisher, Dov Gabbay, and Lluis Vila. 2005. *Handbook of Temporal Reasoning in Artificial Intelligence (Foundations of Artificial Intelligence).* Elsevier Science Inc., New York. 1-24.

September 21: ***Representing Actions, Behaviors, and Effects***
AIMA Chapter 3

Dana Nau, Malik Ghallab, and Paolo Traverso. 2004. *Automated Planning: Theory & Practice*. Morgan Kaufmann Publishers Inc., San Francisco. Sections 2.1-2.3, 2.5, 4.1-4.3.

Presentations: Michael Fisher. 2011. Agent deliberation in an executable temporal framework, *Journal of Applied Logic*, Volume 9, Issue 4, 223-238.

Chitta Baral and Gregory Gelfond. 2011. On representing actions in multi-agent domains. In *Logic programming, knowledge representation, and non-monotonic reasoning*, Marcello Balduccini and Tran Cao Son (Eds.). Springer-Verlag, Berlin. 213-232.

September 26: ***State-space Decision-making***
AIMA Chapter 3 and 5.1-5.3

Dana Nau, Malik Ghallab, and Paolo Traverso. 2004. *Automated Planning: Theory & Practice*. Morgan Kaufmann Publishers Inc., San Francisco. Sections 4.1-4.3, 9.3

September 28: ***Game Tree Search***
AIMA Chapter 3 and 5.1-5.3

Dana Nau, Malik Ghallab, and Paolo Traverso. 2004. *Automated Planning: Theory & Practice*. Morgan Kaufmann Publishers Inc., San Francisco. Sections 4.1-4.3, 9.3

Presentation: Viliam Lisy, Branislav Bosansky, Michal Jakob, and Michal Pechoucek. 2009. Adversarial search with procedural knowledge heuristic. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems*. 899-906.

October 3: ***Representing Uncertainty: Probability Theory Overview***
AMAI Chapter 13

Judea Pearl. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers Inc., San Francisco. Chapter 1.

October 5: ***Probabilistic Reasoning and Graphical Models***
AMAI Chapter 13 and 14.1-14.5

Judea Pearl. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers Inc., San Francisco. Chapter 2.

Presentations: Arthur Choi, Yexiang Xue, and Adnan Darwiche. 2012. Same-decision probability: A confidence measure for threshold-based decisions, *International Journal of Approximate Reasoning*.

Concha Bielza, Manuel Gomez, and Prakash P. Shenoy. 2011. A review of representation issues and modeling challenges with influence diagrams, *Omega*, Volume 39, Issue 3, 227-241.

October 10: ***Bayesian Inference and Reasoning with Bayesian Networks***
AMAI 14.1-14.5

Judea Pearl. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers Inc., San Francisco. Chapter 2.

October 12: ***Bayesian Networks in Dynamic Domains***
AMAI 15.5

Sumit Sanghai, Pedro Domingos, and Daniel Weld. 2005. Relational Dynamic Bayesian Networks, *Journal of Artificial Intelligence Research*, Volume 24, 1-39.

Presentation: Eugene Santos Jr., Deqing Lia, Eunice E. Santosb, and John Korah. 2012. Temporal Bayesian Knowledge Bases–Reasoning about uncertainty with temporal constraints, *Expert Systems with Applications*, Volume 39, Issue 17, 12905-12917.

October 17: ***Planning with Markov Decision Processes (MDPs)***
AMAI 16.5

Dana Nau, Malik Ghallab, and Paolo Traverso. 2004. *Automated Planning: Theory & Practice*. Morgan Kaufmann Publishers Inc., San Francisco. Sections 16.1-16.2.

October 19: ***Uncertainty in MDPs***
AMAI 17.4

Dana Nau, Malik Ghallab, and Paolo Traverso. 2004. *Automated Planning: Theory & Practice*. Morgan Kaufmann Publishers Inc., San Francisco. Section 16.3.

Presentations: Miguel Ramirez and Hector Geffner. 2009. Goal Recognition over POMDPs: Inferring the Intention of a POMDP Agent, In *Proceedings of the Twenty-Second International Joint conference on Artificial Intelligence*. 2009-2014.

Michael Wunder, Michael Kaisers, John Robert Yaros, and Michael Littman. 2011. Using Iterated Reasoning to Predict Opponent Strategies, In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*. 593-600.

Term project proposals due by midnight (11:59:59pm).

October 24: ***Hidden Markov Models (HMMs)***
AMAI 15.3

Lawrence R. Rabiner. 1989. A tutorial on hidden markov models and selected applications in speech recognition, In *Proceedings of the IEEE*. 257-266.

Term project proposal presentations in class.

October 26: ***Reasoning with HMMs***
AMAI 15.3

Lawrence R. Rabiner. 1989. A tutorial on hidden markov models and selected applications in speech recognition, In *Proceedings of the IEEE*. 257-266.

October 31: ***HMMs for Behavioral Modeling***
TBD

Presentation: Xin Liu and Anwitaman Datta. 2012. Modeling Context Aware Dynamic Trust Using Hidden Markov Models. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*. 1938-1944.

November 2: ***MIDTERM EXAM***

November 7: ***Probabilistic Logic***
Theodore Hailperin. 1984. Probability Logic, *Notre Dame Journal of Formal Logic*, Volume 25, Number 3. 198-212.

Nils J. Nilsson. 1986. Probabilistic Logic, *Artificial Intelligence*, Volume 28. 71-87.

Raymond Ng and V.S. Subrahmanian. 1992. Probabilistic Logic Programming, *Information and Computation*, Volume 101, Number 2. 150-201.

November 9: ***Probabilistic Logic and Constraint Satisfaction***
AIMA 6.1

Raymond Ng and V.S. Subrahmanian. 1992. Probabilistic Logic Programming, *Information and Computation*, Volume 101, Number 2. 150-201.

Thomas Lukasiewicz. 1998. Probabilistic Logic Programming, In *Proceedings of the 13th European Conference on Artificial Intelligence.*

K.A. Andersen and J.N. Hooker. 1992. A Linear Programming Framework for Logics of Uncertainty, In *Proceedings of the 26th Hawaii International Conference on Systems Sciences.*

November 14: ***Abductive Reasoning in Probabilistic Logic***
Sheila A. McIlraith. 1998. Logic-Based Abductive Inference, *Technical Report: Stanford University Knowledge Systems Laboratory.*

Gerardo I. Simari and V.S. Subrahmanian. 2010. Abductive Inference in Probabilistic Logic Programs, In *Technical Communications of the 26th International Conference on Logic Programming.*

Presentation: Gerardo I. Simari, John Dickerson, Amy Sliva, and V.S. Subrahmanian. 2012. Parallel Abductive Query Answering in Probabilistic Logic Programs, *ACM Transactions on Computational Logic.*

November 16: ***Utility, Rationality, and Decision-making***
AIMA 16.1-16.5

John Scott. 2000. Rational Choice Theory, In *Understanding Contemporary Society*, Gary Browning, Abigail Halcli, and Frank Webster (Eds.). Sage Publications Ltd. 126-138.

November 21: ***THANKSGIVING HOLIDAY***

November 23: ***THANKSGIVING HOLIDAY***

November 28: ***Classical Game Theory***
TBD

November 30: ***Algorithmic Game Theory***
Eva Tardos and Vijay V. Vazirani. 2007. Computing in Games: Basic So-

lution Concepts and Computational Issues, In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani (Eds.). 3-28.

Joseph Y. Halpern and Rafael Pass. 2011. Algorithmic rationality: adding cost of computation to game theory, *ACM SIGecom Exchange*, Volume 10, Issue 2. 9-15.

Presentation: TBD

December 5: ***Term Project Presentations***

December 7: ***Term Project Presentations***