CSE841 Artificial Intelligence

Dept. of Computer Science and Eng., Michigan State University Fall, 2008

Course web: `http://www.cse.msu.edu/~cse841/`

Description: Graduate survey course in Artificial Intelligence. Types of intelligence, knowledge representation, cognitive models. Goal-based systems, heuristic search, games, deductive systems and expert systems. Computer vision, speech recognition, language understanding, robotics, learning, mental development.

Class: Mondays and Wednesdays, 10.20am - 11:40am, 1225 Engineering Building

Instructor: Juyang (John) Weng, phone: 353-4388, e-mail: `weng@cse.msu.edu`,
home page: `http://www.cse.msu.edu/~weng/`

Office: 3144 Engineering Building

Office hours: Tuesdays and Thursdays 7pm - 8pm or by appointment. Emails and telephone calls are not good for asking questions.

1 Objectives

This is an introduction to Artificial Intelligence for graduate students. It is intended as a survey in the many aspects of AI and will concentrate on the basic concepts of representation and control, followed by some (limited) examples of the use of these concepts in important AI subfields (e.g., Perception, Cognition, Learning, Expert Systems, Robotics). As such, this course will not do much in-depth study of any one area. Rather this course is intended as a good introduction for those who have had no exposure to AI and as a stepping stone for those interested in specific areas of AI. We will try to cover a significant portion of the book during the semester. We will also read a number of papers from the literature to reinforce concepts introduced in class.

Rather than present AI as a loose collection of ideas and techniques, this course will strive to emphasize important unifying themes that occur throughout many areas of AI research (which is also the goal of the text). Further, to take advantages of recent exciting multidisciplinary advances in understanding and modeling the brain, we will link this unifying theme with the brain (which is lacking in the text). These include:

- Many methods can be viewed as context-based searching through a space of learned experience (hypotheses in the text).
- Search should be made more effective through brain's representation of learned knowledge.
- The viewpoint of an intelligent "agent" as an effective one for describing and understanding AI approaches.
- The subject of the brain and its development (not covered by the textbook) potentially will not only unify all the subareas of AI but also lead to wide-front breakthroughs. The realization of this potential needs multidisciplinary research by many researchers who dare to face various challenges of "muddy" tasks.

2 Course Materials

The required text for the course is *Artificial Intelligence* by Russell and Norvig, Prentice-Hall, 2003 (2nd edition). A lot of material covered (e.g. lectures) in the class will be available on the course home page at `http://www.cse.msu.edu/~cse841/`. This web location will also contain links to useful information regarding AI.

Note that the text is primarily a reference for the course, to refresh your memory on basic methods. We will be covering some new material in the course that is not in the text, primarily through the the lectures and assigned papers.

3 Course Prerequisites

This course assumes the following perquisites:

- Programming skills in a high-level language (e.g. C or C++).
- Exposure to concepts in discrete structures, probability/statistics, and algorithmic analysis.

4 Structure of the Course

The following topics will be covered in this course.

1. *Introduction to AI:* What is artificial intelligence? What are some successes and some failures? What are current research topics in AI? What are some of the philosophical issues (chapters 1,26)
2. *Agents:* Agents are computational entities that are embedded in some environment (real world or artificial), and can perceive "states" and can carry out actions. We'll explore some alternative methods for building agents (chapter 2). The concept of agent can serve to unify the rest of the course material.
3. *Search:* One of the fundamental underpinnings of intelligence is the ability to explore alternative courses of perception and actions, either given the model of the environment or autonomously build it. We will cover the basic search methods, as well as heuristic search algorithms (chapters 3,4,6).
4. *Logic:* One powerful method of building intelligent agents is to provide them with *knowledge* of the task. Logic is a very expressive way of formulating general knowledge. We'll study how agents can make inferences given their prior knowledge (chapters 7,8, 9).
5. *Planning:* A hallmark of intelligence is the ability to construct and select among different sequences of actions (plans), without carrying them out. We'll explore how agents can plan using logical models. (chapters 11,12).
6. *Acting under uncertainty:* Environments can often be stochastic, meaning that actions may not have the same effects, and perceptions can be noisy. We'll explore how agents can deal with uncertainty by representing and reasoning about probabilistic knowledge. (chapters 13,14,15,17).
7. *Learning:* Not all the aspects of a programmer designed model can be completely specified by the programmer. No agent can do well without a capability of adapting to changing environments or new environments. We'll explore several popular methods for learning: neural networks, genetic algorithms and reinforcement learning. We'll also take a look at some interesting issues such as "artificial life."(chapters 18,20,21,24)

8. *Development:* Recently, there is a new approach to artificial intelligence which has a potential to unify different AI problems. This new approach is called *Autonomous Mental Development* (AMD) approach or the developmental approach for short. It aims to fully automate the learning process by a program that is not designed for any specific task. This approach is motivated by the human brain and its development from infancy to adulthood. The textbook does not contain this new material and we will add our own.

9. *Conclusion:* Where is AI heading now? What are the current research issues facing the field? (Chapter 27 and more).

4.1 Weekly Course Outline

Explanation of table: Chapters refer to Russell/Norvig text. Readings refer to papers that you are expected to read, and write a short (450 - 500 words) critique. Note the types on the web page. Some are pdf files (readable with acroreader on any system), some are html pages. This is my best guess at the schedule, it is obviously subject to change.

Week	Date	Topics	Text Reference	Reading
1	8/25-8/29	Administrivia, intro to AI	Ch. 1	Turing*
2	9/2-5	Agents	Ch. 2	Boden
3	9/8-12	Search basics	Ch. 3-4	ChineseRoom*
4	9/15-19	Recent advances in search	Ch. 6	McCarthy
5	9/22-26	Logic inference	Ch. 7-8	DeepBlue*
6	9/29 - 10/3	Planning using logic	Ch. 9,11	Mitchell
7	10/6-10	Managing uncertainty	Ch. 13	Muddy Tasks*
8	10/13-17	Prob. reasoning & **Mid-Term** (Wednesday)	Ch. 14	
9	10/20-24	Planning under uncertainty	Ch. 17	Brooks*
10	10/27-10/31	Learning theory	Ch. 18	Zweig
11	11/3-7	Supervised learning	Ch. 20	Kaelbling*
12	11/10-14	Reinforcement learning	Ch. 21	Echo
13	11/17-21	Mental development	Notes	Weng*
14	11/24-26	Perception	Ch. 24	Hofstadter
15	12/1-5	Robotics	Ch. 25	–

The midterm exam is on Wednesday, Oct. 15. The final exam is on Wedsday, Dec 10th, 10:00am - noon

5 Reading

5.1 Doing the work

Every week, you will be required to read a paper all of which is on the class web page. For the ones marked with * in the table above, you are required to turn in a one-page (450 - 500 words) critique of the paper, due Wednesday before the class. Wrong papers will be graded as zero. One page is not a lot of room. Be concise about your analysis. Don't *summarize*, I already have read the paper. *Analyze* it, give me your thoughts.

On Wednesdays when a critique is due, we will discuss the paper de jour near the end of the class.

5.2 List of Papers

The ones marked with * require critique.

1. Week 1: **Turing***: "Computing Machinery and Intelligence"

2. Week 2: **Boden**: "Creativity and Unpredictability"

3. Week 3: **Hauser***: "The Chinese Room Argument"
4. Week 4: **McCarthy**: "Little Thoughts of Thinking Machines"
5. Week 5: **Hsu***: "IBM's Deep Blue Chess Grandmaster Chips"
6. Week 6: **Craven et al.**: "Learning to Extract Symbolic Knowledge from the World Wide Web"
7. Week 7: **Weng***: "Task Muddiness, Intelligence Metrics, and the Necessity of Autonomous Mental Development"
8. Week 9: **Brooks***: "Elephants don't Play Chess"
9. Week 10: **Zweig and Russell**: "Speech Recognition with Dynamic Bayesian Nets".
10. Week 11: **Kaelbling, Littman, Moore***: "Reinforcement learning: A Survey"
11. Week 12: **Hraber, Jones, Forrest**: "The Ecology of Echo"
12. Week 13: **Weng et al.*** "Autonomous Mental Development by Robots and Animals"
13. Week 14: **Chalmers, French, Hofstadter**: "High-level Perception, Representation and Analogy, A Critique of Artificial Intelligence Methodology"

6 Assignments

- Assigned reading. One paper every week. One-page (450- 500 words) critique for each paper marked with *. Points will be deducted for more and fewer words.
- Homework. We will give 6 homeworks. Some require programming. You need to start at least a week early. Every homework is equally weighted.

Week	**Day**	**Date**	**Subject**
2	Wed.	09/03	Critique (Turing) due. Homework 1 handed out
3	Wed.	09/10	Critique (Chinese room) due
4	Wed.	09/17	Homework 1 due. Homework 2 handed out
5	Wed.	09/24	Critique (DeepBlue) due
6	Wed.	10/01	Homework 2 due. Homework 3 handed out
7	Wed.	10/08	Critique (Muddy tasks) due
8	**Mon.**	10/13	Homework 3 due. Homework 4 handed out
9	Wed.	10/22	Critique (Brooks) due
10	Wed.	10/29	Homework 4 due. Homework 5 handed out
11	Wed.	11/05	Critique (Kaelbling) due
12	Wed.	11/12	Homework 5 due. Homework 6 handed out
13	Wed.	11/19	Critique (Weng) due
14	Wed.	11/26	Homework 6 due
15	Wed.	12/03	

A hard copy of each due, including program source program, must be handed in at the beginning (not at the end) of the class on the due date. No late work is accepted. In case of documented schedule crisis (e.g., university team competition) provide formal signed documents.

7 Grading Policy

You will receive grades for this course based on your performance measured as follows.

- Assigned reading – 20%.

  Totally 15 papers and 8 critiques.

- Assigned homework/programming – 25% .

  We will have about 6 homeworks/programming assignments.

- MidTerm – 25% .

  There will be an in-class mid-term exam based on the material covered in the first half of the course. The exam will consist of questions similar to the exercises in the Russell/Norvig textbook.

- Final – 25%.

  There will be a final (cumulative) in-class exam based on the material covered in the whole course.

- Class Participation – 5%.

  This will be based on attendance, and participation in the class discussions.

8 Auxiliary Resources

Here is a source list of reading material pertinent to AI. At various points during the course, you may want more detail on some particular topic. Don't forget to check for AI resources on the web (e.g. the AAAI web page etc.).

8.1 Some AI Books, Supplemental Readings and Journals

The following books are on reserve for this course.

- *Handbook of AI*, Paul Cohen and Ed Feigenbaum (editors), published by Kaufmann. (in reserve collection).
- *Artificial Intelligence*, 3rd edition, by Pat Winston, Addison Wesley.
- *Pattern Classification*, 2nd edition, by Duda, Hart, and Stork, Wiley.
- *Essentials of AI*, Matt Ginsberg, Morgan Kaufmann.
- *An Introduction to Genetic Algorithms*, M. Mitchell, MIT Press. (Supplemental reading for genetic algorithms.)
- *Rethinking Innateness: A Connectionist Perspective on Development*, J.L. Elman and E. A. Bates and M. H. Johnson and A. Karmiloff-Smith and D. Parisi and K. Plunkett, MIT Press. (Supplemental reading for mental development.)
- *Artificial Intelligence journal*, Elsevier press. (A leading AI journal).
- *AMD Newsletters*, the publication of the IEEE CIS AMD Technical Committee. http://www.cse.msu.edu/amdtc/amdnl/
- Conference proceedings: AAAI, IJCAI, WCCI, IJCNN, ICDL.
- *Mental Development Repository* http://www.mentaldev.org/