Introduction to Artificial Intelligence

CSCE 476-876, Spring 2008

URL: `www.cse.unl.edu/~choueiry/S08-476-876`

Berthe Y. Choueiry (Shu-we-ri)
Avery Hall, Room 123B (later, 360) `choueiry@cse.unl.edu`
`Tel: (402)472-5444`

Outline

- Overview of administrative rules
- What is AI?

When do we meet..

- **Lecture**:
  Mon/Wed/Fri, from 11:30 to 12:20 p.m.
  Room 118, Avery Hall
  I come 10 (15?) minutes earlier to answer questions
  and review material from *previous* lectures
  We must leave on time if another class needs to the room.

- **Make-up class/recitation**:
  Wed, from 5:00 p.m. to 6:00 p.m.
  Room 28D Avery Hall

Communications

- Always refer to the syllabus, our contract
- Frequently check the class schedule (web)
  `www.cse.unl.edu/~choueiry/S08-476-876`
- Read your email, regularly
- Contact instructor & TAs:
  cse476@cse.unl.edu
- Broadcast to class (use sparingly):
  cse476-ml@cse.unl.edu

Office hours:

- Instructor:
  Mon/Wed 12:30–01:30 p.m. (right after class) or by appointment
- GTA:
  Jamie Schirf (AVH TBA):
  Office hours: At the Student Resource Center on Tuesday, Thursday 5:00-6:00 p.m. and Friday 10:30-11:30 p.m. or by appointment.
- Professional attitude: respect schedule of TA (& intructor)

Books

- AIMA: **Second edition**.
- Lisp (LWH): **Third edition**.
- Common Lisp the Language (the Steele) **Second edition**.

Topics

1. **Lisp**: fundamental for the course, no way around it
2. Intelligent agents
3. Search
4. Constraint satisfaction
5. Games
6. Logical systems

If time allows:

- Planning systems
- Uncertainty: probability and decision theory

Important warnings

- **Common Lisp is the official programming language. Mandatory.**
- **CSCE 310 is a pre-requisite.**
  If you don't have it, you need to contact the instructor immediately.
- I will come to class 10-15 minutes ahead of schedule to answer questions.
- Beyond office hours, communicate with us by email as much as possible.
- Class time is limited. **Do your required reading**.

Related courses at CSE:

- Artificial Intelligence (976)
- Constraint Processing (421/821 & CSE990-06)
- Neural Networks & Genetic Algorithms (479/879, 974, 976)
- Machine Learning (478/878)
- Multiagent Systems (475/896)
- Logic (465/865)
- Database (413/813, 913, 914)

Course load

- Required and recommended reading: AIMA & LWH
- Homework: Programming (CL), theoretical, library-search
  To be submitted **before** class, late-return policy, indicate effort
- (Surprise) Quizzes: frequent, cover class discussions & required reading, cannot be made up
- Tests: Pretest, midterm, and final
  Fixed schedule: cannot be taken in advance or made up
  General policy: closed books, crib-sheet policy

Student's responsibility

- Account on `cse` (or csnt), using `xemacs` and `lisp`
- No plagiarism, heavily sanctioned. Review policy of CSE
- Always acknowledge sources, help, individuals, url, etc.
- Attendance not mandatory, however students are responsible for material covered and quizzes taken
- Professional behavior:
  don't miss classes, don't come late to classes, don't solicit help beyond office hours without an appointment

Grading policy

- Homework 30%
- Pretest 5%
- Quizzes 15%
- Midterm 25%
- Final 25%

Secure a good grade

- Bonus for full attendance
- Glossary: weekly & final. Alphabetically sorted. (Up to 8%)
- Bonus for solving occasional riddles
- Bonus for finding errors of the instructor

How well you are doing: feedback mechanisms

- Quizzes are corrected in class.
- Homework and glossaries are promptly corrected.
- Grades are listed in great details at board outside my office.
- You have 5 working days to claim grade adjustment. Strictly reinforced.
- Students who are not performing are contacted directly. Grades are monitored, but I cannot force you to work.
- Your suggestions for improving the course and our feedback mechanisms are *most welcome*, carefully considered, and implemented as quickly as possible.
- Please let us know what other feedback you expect.

Other resources

- Books on reserve at the Math Library (Avery)
- LL collection, dictionaries, and reference books
- On-line pointers to AI, Lisp, etc. (course and schedule pages)
- Student's catch from the web

Pretest

- Scheduled for Friday, Jan 18, 2008
- One part to be completed in the class: crib sheet policy
- One part to be completed at home: collaboration, discussion strictly forbidden
- Content: basic knowledge of mathematics, logic, algorithm, data structure, complexity

Goal of AI

- Understand intelligent entities (reasoning mechanisms)
- Build intelligent entities (systems)
  contrast with cognitive science and philosophy

→ Build computers with human-level intelligence.. or better
(human reasoning exhibits systematic errors)

Using: slow, tiny brain, biological or electronic

In order to: perceive, understand, predict and manipulate a far more complex world

Proof of feasibility: human beings
just look in the mirror :–)

New discipline, old topic

AI is a new discipline (vs. physics):
- term coined in 1956 by John McCarthy
- task is enormous, opportunities are wide, easy to make a difference
- Einstein is (probably) yet to come

Study of Intelligence is an old topic. Philosophy: learned but speculative

Advent of computers introduced a new experimental and theoretical discipline: theories can now be tested
⟶ out of the armchair, into the fire

Early Systems were naive (rule-based, etc.)
Paradigms are getting more difficult, elaborate, richer, more subtle

Focus and fields

General:

- perception
- logical reasoning

Specific: (task oriented)

- chess
- proving mathematical theorems
- pun writing
- diagnosing diseases
- planning/scheduling tasks of building construction

A truly universal field

Often scientists/engineers become AI researchers:
want to formalize, systematize, automate the intellectual tasks they are trained to carry out (electrical engineers, civil engineers, medical doctors)

Sometimes, AI researchers delve into specific fields to apply their methods (biology, power systems)

"The exciting new effort to make computers think ... *machines with minds*, in the full and literal sense" (Haugeland, 1985) "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning ..." (Bellman, 1978)	"The study of mental faculties through the use of computational models" (Charniak and McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act" (Winston, 1992)
"The art of creating machines that perform functions that require intelligence when performed by people" (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991)	"A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes" (Schalkoff, 1990) "The branch of computer science that is concerned with the automation of intelligent behavior" (Luger and Stubblefield, 1993)

Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

Dimensions for classification:

Vertical: concern, focus of efforts
→ thought process and reasoning
→ behavior and action

Horizontal: evaluation of success
→ against human performance
→ against ideal concepts of intelligence

Rationality = do the right thing

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

Classification contrast human & rationality:
- Human: empirical science, hypothesis and experimental
- Rationality: mathematics + engineering

No right/wrong, all four approaches are valuable