welcome to cis32 = artificial intelligence (AI)

MW 12.15 - 1.30pm 214 New Ingersoll

Professor Elizabeth Sklar

email: *sklar@sci.brooklyn.cuny.edu*
web: *http://www.sci.brooklyn.cuny.edu/~sklar*
office: 1417 Ingersoll
office hours: mondays 11am-12noon

class web page:
http://www.sci.brooklyn.cuny.edu/∼sklar/ai

course objectives

- To introduce you to some of the basic theory and practical techniques in artificial intelligence.
- In particular, this course will teach you about:
  - Introduction to Artificial Intelligence, Autonomous Agents
  - Problem solving, Search, Heuristic methods
  - Introduction to Robotics
  - Expert Systems, Intelligent Tutoring Systems
  - Neural Networks, Evolutionary Computation, Artificial Life
  - State space Learning, Game Playing, Knowledge Representation, Uncertainty
  - Propositional Logic, Predicate Logic, Logic-based Agents
  - Planning, Partial-order planning, Decision-theoretic planning
- You will also get to play with robots – *Hey that's the best part :-)*

your input

- To give you experience of applying both the theory and practical techniques.
- In other words there are:
  - homeworks (30%)
  - projects (30%)
  - exams (40%)
- Homeworks will give you the chance to use the theory.
- Projects will give you the chance to use the practical techniques.

resources

- lectures
- handouts, primarily from three sources:
  - Artificial Intelligence: A Modern Approach (2nd edition)
    by Stuart Russell and Peter Norvig, published by: Prentice Hall (2003)
  - Artificial Intelligence: A New Synthesis
    by Nils J. Nilsson, published by: Morgan Kaufmann (1998)
  - An Introduction to Multiagent Systems
    by Michael Wooldridge, published by: John Wiley & Sons (2002)
- lecture notes
- class web page
  - *http://www.sci.brooklyn.cuny.edu/∼sklar/ai*

office hours

- I will hold office hours before class on Mondays 11am-12noon
- I will also hold *virtual office hours*
  AIM screen name: $agentprof$

assessment

- out of 100 points

30	points	=	6 homework assignments (5 pts each)
30	points	=	2 projects (15 points each)
40	points	=	2 exams (20 pts each)
100	points		

- *note that the midterm date is tentatively set for* March 27

a word about homeworks

- Should be done on your own, as much as possible
- Get help from me, friends
  but you must acknowledge all help received by citing the names of those who helped you.
- This not only protects you from being accused of cheating, but also protects you in case your helper gives you misinformation

homeworks: submission policy

- Homeworks are due on the day that they are due.
- Here are the rules — please know them well:
  1. Electronic submissions are due at 6am on their due date
     OR
  2. Hardcopies must be brought to class on the due date and deposited in the homework box at the front of the classroom.
  3. If your hardcopy does not make it into the box, it will not be accepted and you will get 0 for the homework.
  4. If you must miss class, have a friend deposit your hardcopy.

to every rule, there is an exception...

- Exceptions and extensions are possible, primarily based on MEDICAL EMERGENCIES.
- Circumstances must be documented and suitable arrangements will be made.
- You must consult me via email on an individual basis.
- You must consult me BEFORE you need the exception/extension, not afterwards.
- Submission details for projects will be posted on the class web page.

regrade policy

- If you feel that there was an error in grading your homework, project or exam, then you need to write on a piece of paper a description of the error.
- STAPLE the paper to your homework, project or exam and leave it with me to be regraded.
- Know that I mark with a list of expectations for each homework assignment, project and exam problem, knowing where to take off points — so if your complaint is that too many points were taken off for one kind of mistake or another in your program, then generally those types of things will not change in a regrade.

regrade policy, continued

- If there is a genuine error in the marking, like I thought something was missing, but it is really there, then you will likely get points restored.
- HOWEVER, a regrade means that the entire assignment or exam will be remarked, so be aware that your mark can go DOWN as well as up.
- Regrades take while to process, so be patient — if you need the work to study from, then make a copy of it before you turn it in for a regrade.

a word to the wise

You all know that:

- You should save early and save often!
- Disk drives crash.
- Floppies have bad sectors.
- Power supplies fail.
- Monitors die, and
- Paper print-outs are the best form of backup storage known to mankind.

So, you'll know that problems resulting from ignoring the above are not acceptable excuses for late submission of projects or homework.

a word about lectures

- Brief lecture notes will be placed on the web page after every lecture.
- These will be linked to the syllabus page.
- But they are NOT A SUBSTITUTE FOR COMING TO CLASS.
- I know, I used to skip classes too.
- If you must miss a class, YOU are responsible for getting notes from someone who did come to class

a word about note-taking

- I will try to post lecture notes on the web before class BUT:
  - you learn better when you actually have to write things down yourself.
  - just reading along with my notes makes you sleepy.
  - everything I say is NOT in the lecture notes, but anything I say MIGHT be on an exam or in a homework, so you need to take notes on what I say
  - sometimes there are mistakes in the lecture notes which get caught during class; so you will only get the correct version if you come to class and take notes.

a word about exams

Exams are not a great way of assessing what people know, but they:

- Are the only way I know you are doing your own work.
- Are the only way YOU know you are doing your own work.
- Are not hard if you really know the material.
- Notice my weighting scheme for exams.
  1. midterm: 20%
  2. final exam: 20%

a word about feedback

- Homeworks, projects and exams let me know how you are doing.
- In a way, they let me know how I am doing, as a reflection of how you are doing.
- But, I welcome feedback from you:
  - email,
  - anonymous written notes,
  - instant messages,
  - comments during office hours.

a word about academic integrity

- The work you submit for assessment should be completed ON YOUR OWN.
- You may get help from me, friends.
- You must acknowledge all help given.
- You must not download material from the web and submit it as your own work.
- You must not mail code or copy files.
- If someone asks you to do this, *JUST SAY NO!*.

topics covered

I plan to cover the following:

- Introduction to Artificial Intelligence, Autonomous Agents
- Problem solving, Search, Heuristic methods
- Introduction to Robotics
- Expert Systems, Intelligent Tutoring Systems
- Neural Networks, Evolutionary Computation, Artificial Life
- State space Learning, Game Playing, Knowledge Representation, Uncertainty
- Propositional Logic, Predicate Logic, Logic-based Agents
- Planning, Partial-order planning, Decision-theoretic planning

Not necessarily in that order...
For full details see the class syllabus page.

So, what is AI?

- AI is both science and engineering:
  - the *science* of understanding intelligent entities — of developing theories which attempt to explain and predict the nature of such entities;
  - the *engineering* of intelligent entities.
- four views of AI
  1. AI as *acting humanly*
     - as typified by the Turing test
  2. AI as *thinking humanly*
     - cognitive science.
  3. AI as *thinking rationally*
     - as typified by logical approaches.
  4. AI as *acting rationally*
     - the intelligent agent approach.

acting humanly

- A problem that has greatly troubled AI researchers: *when can we count a machine as being intelligent*?
- Most famous response due to Alan Turing, British mathematician and computing pioneer:

  Human interrogates entity via teletype for 5 minutes. If, after 5 minutes, human cannot tell whether entity is human or machine, then the entity must be counted as intelligent.

- No program has yet passed Turing test! (Annual Loebner competition & prize.)
- A program that succeeded would need to be capable of:
  - natural language understanding & generation;
  - knowledge representation;
  - learning;
  - automated reasoning.
- Note no *visual* or *aural* component to basic Turing test — augmented test involves video & audio feed to entity.

thinking humanly

- Try to understand how the mind works — how do we think?
- Two possible routes to find answers:
  - by *introspection* — we figure it out ourselves!
  - by *experiment* — draw upon techniques of psychology to conduct controlled experiments. ("Rat in a box"!)
- The discipline of *cognitive science*: particularly influential in *vision*, *natural language processing*, and *learning*.

thinking rationally

- Trying to understand how we *actually* think is one route to AI — but how about how we *should* think.
- Use *logic* to capture the *laws of rational thought* as *symbols*.
- *Reasoning* involves shifting symbols according to well-defined rules (like algebra).
- Result is *idealised* reasoning.
- Logicist approach theoretically attractive.

- Lots of problems:
  - *transduction* — how to map the environment to symbolic representation;
  - *representation* — how to represent real world phenomena (time, space, ...) symbolically;
  - *reasoning* — how to do symbolic manipulation *tractability* — so it can be done by real computers!
- We are still a long way from solving these problems.
- In general logic-based approaches are unpopular in AI at the moment.

acting rationally

- Acting rationally = acting to achieve one's goals, given one's beliefs.
- An *agent* is a system that perceives and acts; intelligent agent is one that acts rationally w.r.t. the goals we delegate to it.
- Emphasis shifts from designing *theoretically best* decision making procedure to best decision making procedure possible in circumstances.
- Logic may be used in the service of finding the best action — not an end in itself.

- Achieving *perfect* rationality — making the best decision theoretically possible — is not usually possible, due to *limited resources*:
  - limited time;
  - limited computational power;
  - limited memory;
  - limited or uncertain information about environment.
- The trick is to *do the best with what you've got*!
- This is easier than doing perfectly, but still tough.

about me

- undergrad: Barnard College, Columbia University, Computer Science / English
- grad school: Brandeis University, Computer Science, MS, PhD
- previous work experience: MIT / Lincoln Laboratory, GTE Government Systems, consulting
- previous teaching: Monash University, Australia; Boston College, USA; Columbia University, USA
- research interests:
  - interactive learning systems
  - multiagent simulation of learning environments
  - educational robotics

about you

- Please take out a piece of paper and write down...
  1. Your name.
  2. Your email address.
  3. Your class and major OR if you are a non-matriculating student, categorize yourself.
  4. Why you are taking this course?
  5. What you hope to learn in this course?
  6. What do you think artificial intelligence is?
  7. Tell me one fun thing you did over the break.
- ...and give it to me before you leave.