Syllabus

Spring 2016

Course:	CMSC 389: Artificial Intelligence
Time:	TR 12:30pm
Room:	Ruffner 352
Website:	`http://cs.longwood.edu/courses/cmsc389/`

A programming intensive course covering the theory and techniques of artificial intelligence (AI) with implementations using both statistical and nonstatistical AI methods. Covered topics will span the central AI problems of planning, learning, and reasoning. Prerequisite: CMSC 262. 3 credits.

Professor:	Don Blaheta
Office:	Ruffner 337
Phone:	x2191
Email:	`blahetadp@blahedo.org` [1]
Office hours:	Tue 2–3:30pm; Wed 4–5:30pm; Fri 2–3pm; and in the Hardy House Mon 9–11pm

General info

Artificial intelligence is in many ways a moving target. Once a problem is solved, or at least once its difficulties are somewhat understood, it is frequently no longer considered AI! Nevertheless, there are a few key areas that remain central to the idea of intelligence, and that feature heavily in AI textbooks. In this course, we will focus on three of them: problem space search, statistical reasoning, and neural networks. By the end of the course, you'll be expected to know several of the main algorithms and frameworks for reasoning and learning, but more importantly, you'll be expected to understand what makes them relevant, why a researcher might choose them, and where their strengths and weaknesses lie.

The book for this course will be Russell and Norvig's *Artificial intelligence: a modern approach*, 3rd edition. ISBN 978-0-13-604259-4.

[1] Or `blahetadp@longwood.edu` if you'd rather, but I prefer the off-campus one and check it more frequently.

Course objectives

At the end of this course, the successful student will be able to:

1. analyse a range of different types of problems in terms of problem spaces, and frame their solutions in terms of searching those spaces;
2. induce statistical models based on empirical data and use them to make predictions;
3. implement a standard learning algorithm to construct an abstract model based on a training set.

Calendar

Though there are exceptions, as a general rule Tuesdays will have less reading assigned, as that will be when programming work is due; and Thursdays will be when heavier reading is assigned and more theoretical homework is due. If the reading looks too heavy to accomplish in two days (or you're busy that week), you're always welcome to start it early.

The projects and exams will follow roughly this schedule:

	Out	Checkpoint	Due
Project 0 (word ladders)	19 Jan	26 Jan	2 Feb
Project 1 (game AI)	9 Feb	16 Feb	1 Mar
Exam 1	1 Mar		4 Mar
Project 2 (swype)	15 Mar	22 Mar	5 Apr
Project 3 (OCR)	5 Apr	12 Apr	26 Apr
Exam 2	29 Apr		6 May

but note that this plan may be adjusted or adapted in case of weather or general deadline slippage.

Topics and reading assignments (tentative)

Wk	T	R	
	January		
1	**19** Introductions, Administrivia Algorithm design for AI problems ***Proj 0 out***	**21** Ch. 2 Agents and environments	
2	**26** * § 3.1 Formally specifying problems and problem spaces ***Proj 0 checkpoint***	**28** §§ 3.2–3.5 More formal specification A* revisited	
	February		
3	**2** § 3.6 Evaluating and improving heuristics Admissibility and informedness ***Proj 0 due***	**4** §§ 5.1–5.3 Minimax Game day	
4	**9** — Representing game states, moves Evaluating heuristics, endgame ***Proj 1 out***	**11** §§ 5.4–5.5 Real time; stochastic evaluation; Partially observable games	
5	**16** — Time-constrained AI ***Proj 1 checkpoint***	**18** §§ 6.1–6.3 Constraint-satisfaction Backtracking	
6	**23** §§ 6.4–6.5 Search and general problem-solving	**25** §§ 7.1–7.5 Propositional logic Theorem proving	**
	March		
7	**1** — Math proofs ***Proj 1 due*** ***Exam 1 out***	[**Exam 1** due Fri 4th **no class**]	

SPRING BREAK

* **26 January**: Deadline to add/drop classes (5pm)
** **26 February**: Deadline to elect pass/fail option (5pm)

Wk		T	R
		March	
8	*	**15** Ch. 13 Probability and Bayes' Law Data-driven algorithm design ***Proj 2 out***	**17** §§ 14.1–14.2 Bayes nets
9		**22** §§ 15.1–15.3 Markov models Hidden Markov models ***Proj 2 checkpoint***	**24** §§ 22.1–22.3 Statistical NLP
10		**29** TBA HMMs in NLP Noisy channel model	**31** §§ 18.1–18.2, 18.7 Neural networks
		April	
11		**5** — Multilayer neural networks Backpropagation ***Proj 2 due*** ***Proj 3 out***	**7** §§ 18.3–18.4, 18.6 Decision trees Linear models Evaluating models
12		**12** § 18.5 Theory of learning ***Proj 3 checkpoint***	**14** TBA Deep Dream Feature selection
13		**19** §§ 10.1–10.2 Planning as search	**21** §§ 10.3–10.5, 11.1 Planning algorithms
14		**26** TBA Responsive agents, emergent systems ***Proj 3 due***	**29** — Planning: monolithic systems vs emergent behaviour ***Exam 2 out***
		May	
			Exam 2 due **Fri, 6 May @10:30am**

* **14 March**: Deadline to withdraw from a class (5pm)

Graded work

Projects. Much of this course will revolve around the programming projects: one warm-up project (worth 10%) and three full-size projects (worth 15% each; best score counts for 20%). They'll be challenging and somewhat open-ended, but they're the best way to achieve real understanding of this material. They are "collaborative": you're encouraged to discuss them and bounce ideas off each other, although in the end you have to write your own program yourself. — 60% collaborative

Homework. Most weeks, usually on Tuesday, I will assign a short written homework at the end of class to be due at the beginning of the next. Each homework will proceed in two rounds: in response to your first handin, I'll give feedback (but no grade); after you have revised it, I'll assign a grade. Each problem will get 5, 3, or 0 points. The homeworks are group work: you can work with anyone in the class (or on your own if you prefer), and mark the names of the whole group at the top of a single handin. These will make up 15% of the grade. — 15% group

Exams. There will be two exams, one at midterm and one for the final. Both will be take-home, and you will be given several days to work on them. They are non-collaborative: you are not permitted to discuss the exam with anyone else other than me. Each exam is worth 10% of the final grade. — 20% non-collaborative

Preparedness. This is a catch-all category for things I do to keep you honest about being prepared for class: reading quizzes, whiteboard work, in-class project updates, and so on. — 5%

Grading scale

I tend to grade hard on individual assignments, but compensate for this in the final grades. The grading scale will be approximately as follows:

A−	[85, 90)	A	[90, 100)	[2]	
B−	[70, 75)	B	[75, 80)	B+	[80, 85)
C−	[55, 60)	C	[60, 65)	C+	[65, 70)
D−	[40, 45)	D	[45, 50)	D+	[50, 55)

While there will be no "curve" in the statistical sense, I may slightly adjust the scale at the end of the term if it turns out some of the assignments were too difficult.

[2] Alas, no A+, unfortunately.

Policies

Honor code policy

Above all, I ask and expect that you will conduct yourself with honesty and integrity—and not to ignore the other ten points of the Honor Code, either. Take pride in what you are capable of, and have the humility to give credit where it is due.

The two main forms of academic dishonesty are "cheating" and "plagiarism". "Cheating" is getting help from someplace you shouldn't, and "plagiarism" is presenting someone else's idea as if it's your own. If you ever find yourself inclined towards either of these, know that there are always other, better options. Persevere! See my website[3] for some discussion and examples of how to steer clear of these problems, and feel free to come talk to me if you need help finding some of those other options (even if it's for another course).

Cheating or plagiarism (on any assignment) will normally receive a *minimum* penalty of a lowered *course* grade, ranging up to an F in the course. Cases will also be turned in to the Honor Board. But: I believe in your potential, and I hope that you will, or will grow to, observe this policy not simply to evade punishment but positively as a matter of character.

Systems and environments

In general, for the programming problems, you'll be free to choose whatever programming language you think is most appropriate to the problem. The main supported languages are Scheme/Racket, Java, Python, C, and C++, though others are possible; the main supported systems are those in the Hardy House lab, but if you want to use your own computer you're welcome to do so (but *you* are responsible for making sure your program runs on our systems before you hand it in).

Accommodations

If you have any special need that I can accommodate, I'm happy to do so; come speak to me early in the term so we can set things up. If you have a documented disability, you should also contact Longwood's Office of Disability

[3] `http://cs.longwood.edu/~dblaheta/collab.html`

Resources (Graham Hall, x2391) to discuss some of the support the college can offer you. All such conversations are confidential.

Attendance and late policy

Attendance is required, and assignments must be turned in on time. That said, if you have a good reason to miss class or hand something in late, I tend to be fairly liberal with extensions if you ask in advance. (Good reasons do include assignments due for other classes.) (And medical and family emergencies are exempted from the "in advance" part, of course. But contact me ASAP.)

Frequent absence will result in a lowered participation grade; habitual absence may in extreme cases result in a failing grade for the class. *Unexcused* late assignments will normally be given a zero.

Inclement weather policy

I don't plan to cancel class for weather unless the entire college shuts down. If you are commuting or are otherwise significantly affected by a weather event, use your own best judgement; and if you do miss class for this reason, contact me as soon as possible to make up missed work.

Early bird policy

Nobody's perfect, and on occasion an assignment gets written a little unclearly (or, once in a while, with an actual error in it). If you catch one and bring it to my attention early, so that I can issue a clarification or correction to the rest of the class, there'll be some extra credit in it for you.