Syllabus

Spring 2022

Course:	CMSC 389: Artificial Intelligence
Time:	TR 12:30pm
Room:	Stevens 118
Websites:	`https://canvas.longwood.edu/courses/1304686`
	`https://www.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:	Rotunda 334
Phone:	x2191
Email:	`blahetadp@longwood.edu`
100% office hours:	Mondays 2–3pm; Tuesdays 11am–12:30pm; Wednesdays 3-4pm; Thursdays 2–3:30pm

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.

Covid-19 notes

There are a number of policies specific to running a class in a pandemic that I wanted to put early in the syllabus to get your attention.

Attending class. There are two ways you can attend class: in person, or via Zoom link. Either mode of attendance is equivalent for purposes of evaluating your presence and participation; if you attend via Zoom link,

- you must have a reason, and
- you must say what it is,

but I don't need any medical detail and if it's not directly covid-related I'm not going to police that. (Basically: be an adult and make good choices.) Unlike last year, I can't promise that the Zoom experience will be anywhere near equivalent to the in-person experience; in fact, it almost certainly won't be. But if you are quarantined, or otherwise just can't attend in person on a particular day, zooming is better than total absence.

Important note: I plan to, in general, turn on Zoom every day; but it's possible I won't if I don't know for sure someone will be attending that way. Try to fill out the "why am I zooming today" link as soon as you know you'll need to.

Medical needs. There are a number of medical reasons why attending class in person may not be appropriate for you. Obviously, if you receive a positive Covid-19 test, you will need to remain in isolation and attend class via Zoom link. Even without a positive confirmation, if you are feeling some symptoms or have been exposed or are awaiting test results, attending via Zoom link from quarantine is most appropriate (*especially* if you are unvaccinated).

> **More serious medical needs.** If you are feeling serious symptoms of Covid-19 (or some other sickness), your priority should be on dealing with that. If you end up missing class sessions and/or assignments due to being sick, notify me when you can and then let me know when you're on the upswing so we can plan out how to get you caught up.

What if the professor gets sick? Same as for students: if I'm feeling sick, I'll zoom myself into the class. If necessary I can teach from a zoom window on the projector screen; I'll send an email with instructions as soon as I know I need to do this.

Wearing a mask. As of the start of the semester, all students attending class in person must be wearing a mask or other appropriate face covering. Coverings that are acceptable include some kinds of folded bandannas, gaiters,

or scarves, as long as: it covers both your nose and mouth, with two layers of cloth, fitted relatively snugly around the edges, and reduces aerosols (i.e. it's relatively tightly woven, not very stretchy, and doesn't have an "exhaust port"). If you show up to class without a face covering, you will be required to put one on or leave. If you are medically unable to wear a mask, please contact the Accessibility Resources Office to help you work out an accommodation.

It is possible that the university's masking rules will relax at some point. Even if/when they are not generally required, I will encourage you to wear a mask if you are unvaccinated, if you are mildly symptomatic (e.g. "it's probably allergies but just in case"), or if it makes you feel safer or more comfortable to do so.

All-online? It's still slightly possible that at some point in the semester we'll have to move all-online to handle an outbreak (or that *I* will land in quarantine). Should that happen, we will migrate the course to Zoom meetings but otherwise carry on. I expect that this course will remain largely synchronous (i.e. we meet at our regular class time) even if we go remote, but some calendar dates may be adjusted.

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.

Resources

In general I will expect that you've read about each day's topic in advance of coming to class for the day. Two textbooks are available:

- Lucci and Kopec, *Artificial intelligence in the 21st century*, 2nd edition. ISBN 978-1-94227000-3.
- Russell and Norvig, *Artificial intelligence: a modern approach*, 3rd edition. ISBN 978-0-13-604259-4.

You are welcome to get your own copy of either one (Russell and Norvig is more dense but very dry; Lucci and Kopec are better at working through examples) BUT I have put my copies on semi-permanent loan to the ACL lounge library and you can read them there.

There will also sometimes be links to Wikipedia or other sites, and you can also seek your own references on the posted topics.

Time

I figure that I have on average about 9 hours of your time every week, including class time as well as reading, practice, and projects. If you find you're regularly spending substantially more time than this, please do come discuss it with me, so that we can ensure you're making the most effective use of your time.

Calendar

The projects and exams will follow roughly this schedule:

	Out	Checkpoint	Due
Project 0 (word ladders)	13 Jan	20 Jan	1 Feb
Project 1 (game AI)	1 Feb	8 Feb	22 Feb
Exam 1	1 Mar		4 Mar
Project 2 (autocorrect)	15 Mar	22 Mar	5 Apr
Project 3 (OCR)	5 Apr	12 Apr	26 Apr
Exam 2	28 Apr		3 May

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

Topics (see website for reading assignments)

Wk	T	R
	January	
1		**13** Introductions, Administrivia Agents and environments ***Proj 0 out***
2	**18** Formally specifying problems and problem spaces Informed search	**20** * Optimal search Branch and bound A* revisited Admissibility and informedness ***Proj 0 checkpoint***
3	**25** Search, continued	**27** Minimax revisited Alpha-beta pruning
	February	
4	**1** Game day ***Proj 0 due*** ***Proj 1 out***	**3** Representing game states, moves Evaluating and improving heuristics Real time/time-constrained AI Stochastic, partially observable games
5	**8** Project 1 implementation design ***Proj 1 checkpoint***	**10** Fuzzy logic
6	**15** Fuzzy logic cont'd Bayesian reasoning revisited	**17** ** Decision trees, Linear models Entropy
7	**22** (cont'd) ***Proj 1 due***	**24** Propositional logic Theorem proving Math proofs
	March	
8	**1** Predicate logic Unification ***Exam 1 out***	[***Exam 1*** due Fri 4th **no class**]

* **20 January**: Deadline to add/drop classes (5pm)

** **18 February**: Deadline to elect pass/fail option (5pm)

Wk	T		R
	March		
	SPRING BREAK		
9	**15** Hidden Markov models Noisy channel model ***Proj 2 out***		**17** (cont'd)
10	**22** Supervised learning Training and testing Evaluating models ***Proj 2 checkpoint***		**24** Genetic algorithms
11	**29** Genetic algorithms cont'd	*	**31** Neural networks Perceptron models
	April		
12	**5** Multilayer neural networks Backpropagation ***Proj 2 due*** ***Proj 3 out***		**7** Backpropagation cont'd Genetic-neural programming
13	**12** Planning problems ***Proj 3 checkpoint***		**14** Planning as search
14	[Research Day **no class**]		**21** Planning algorithms Production systems Intelligent agents
15	**26** Responsive agents, emergent systems ***Proj 3 due***		**28** Planning: monolithic systems vs emergent behaviour ***Exam 2 out***
	May		
	Exam 2 due **Tue, 3 May @2pm**		

* **30 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 10% 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. — 50% collaborative

Each will have a checkpoint that I'll describe in detail elsewhere, plus a final handin. The checkpoint and the final versions are due at 4pm on their respective due dates unless I say otherwise.

Collected homework. From time to time, I will assign a written homework at the end of one class to be due at the beginning of the next. Each homework will proceed in two rounds: first, I give feedback (but no grade); on a revision, I assign a grade (but no further feedback). Each problem gets 10, 6, or 0 points. Homeworks are group work: you can work with anyone in the class, or on your own if you prefer, and hand in one copy for the group. — 15% group

These will make up 15% of the grade. All homeworks are due at the beginning of class on the due date.

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 15% of the final grade. — 30% non-collaborative

Engagement. You need to be actively engaged in this class. Engagement comes in many forms, but I expect that you will be interacting with your classmates, and with me, both in class (in-person or Zoom) and in the Slack channel. General engagement will be evaluated in two-week blocks—so you don't need to artificially say a thing every day—and it's ok if most of your engagement is via Slack as long as *some* of it is in class (spoken or in the Zoom chat). In addition, there will be occasional required interactions via Canvas that will be considered part of the engagement grade (especially at the start of the semester and if we lock down at some point).

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, 95)	A+	[95, 100]
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. Final grades of A+ are recorded as an A in the grading system. Final grades below the minimum for D− are recorded as an F.

Note that *individual* grades recorded in Canvas should be accurate (and you should let me know if there's a data entry error!), but *averages* as computed by Canvas sometimes are not, if the averaging is complex or (especially) if an individual student has a special case scenario. The reference gradebook is my own spreadsheet, and while I will try to make Canvas reflect it (including averages) as well as I can, Canvas can't always handle it.

Policies

You can find several university-wide course policies at
`http://www.longwood.edu/academicaffairs/syllabus-statements/` .

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[1] for some discussion and examples of how

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

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 lowering the *course* grade by a full letter, and may range at my discretion 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 Advanced Computing 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 Accessibility Resources Office (Brock 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; and if the campus closes, I'm likely to hold class in some form by zoom instead (check your email). If you are commuting or are otherwise significantly affected by a weather event, use your own best judgement (and remember that zoom is an option); and if you do miss class for this reason (e.g.: power's out too), 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.