Syllabus
CMSC 389: Artificial Intelligence

Spring 2026

Time:	MW 4pm
Room:	Rotunda 356
Website:	`https://cs.longwood.edu/courses/cmsc389/`

A programming intensive course covering the theory and techniques of classic 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 3–4pm; Tuesdays 1–2pm; Wednesdays 10–11am; Thursdays 11am–noon

General info

What a time to be stuying artificial intelligence! Since 2022, "AI" has been all over the media as the biggest newest thing, particularly in the form of text-driven large-language-model generative AI (like ChatGPT). The overall field is much older than that, one of the oldest subfields of computer science, and in this course we will investigate several classic areas of artificial intelligence (some of which directly pertain to the big genAI systems, and some of which appear elsewhere in the AI ecosystem). 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.

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.

AI Policy

My general feeling about AI is this: AI is a tool. Use it when it's helpful, don't use it when you could do it better or faster yourself.

That said, there are certain skills that programmers and computer scientists will need to develop and execute without the help of AI, slightly because AI might not *always* be available but mostly because you'll need to be able to evaluate and debug the code that the AI (or other programmers) have given

you. Thus for assignments that are about *developing* your programming skills (labs, homeworks, projects), I'm going to discourage use of AI until you've given a few solid attempts without. For assignments that are *assessing* your skills (exams) I'll have specific instructions on whether you are or are not allowed to use generative AI to assist. *In general* tasks that you're doing on your own time will permit use of AI, but please attend to specific instructions on each assignment.

I will expect that when you *do* use generative AI, you will document it: say which AI system you used and what help it gave you. **In a comment or embedded link, you should include the "share" URL that lets others view your prompts along with the AI's responses.** Some assignments will have additional instructions how to document this.

(Note that although Longwood's Honor Code does not inherently ban the use of AI, some other professors seem to think it does, so for your safety you should check with each professor before using it in their class.)

Attendance, absence, and zoom (fka "Covid-19 notes")

I have a few policies that originally evolved in response to the pandemic but I've decided they're just good policy so I kept them. Here's the gist: It's really easy to keep zoom open for every class, and it's not nearly as good as in-person attendance but way better than total absence. So I open Zoom every day and ask you to make good choices.

Attending class. There are two ways you can attend class: in person, or via Zoom. Either mode of attendance is sufficient to be "attending" (but not necessarily engaging or participating). If you attend via Zoom,

- 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 illness-related I'm not going to police that. (Again: be an adult and make good choices.) The Zoom experience is nowhere near equivalent to the in-person experience and is not a replacement for it; but if you are quarantined, or otherwise just can't attend in person on a particular day, zooming is better than total absence.

Zooming vs masking. Although we've moved to "endemic" on Covid-19, remember that masking is still a tool in our toolkit: if you're feeling a bit

sniffly, you can still wear a mask. We all have masks, we all got really good at wearing them, and it's a courtesy to your classmates to take this easy step to decrease the likelihood of spreading anything. (Including colds and other stuff! Masks help us not spread *lots* of things.)

What if the professor gets sick? Same as for students: if I'm feeling a little sniffly, I'll wear a mask, and if I am more seriously sick (but well enough to teach), I'll zoom myself into the class. If necessary I can teach from a zoom window on the projector screen (and have done so!); I'll post to the Slack and send an email with instructions as soon as I know I need to do this.

Time

I figure that I have on average about 9 hours of your time every week, including class time as well as reading, practice, homework, 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	Followup
Project 0 (word ladders)	14 Jan	21 Jan	4 Feb	18 Feb
Project 1 (game AI)	4 Feb	11 Feb	25 Feb	4 Mar
Exam 1	2 Mar		4 Mar	
Project 2 (autocorrect)	16 Mar	23 Mar	6 Apr	13 Apr
Project 3 (tbd)	6 Apr	13 Apr	27 Apr	8 May
Exam 2	29 Apr		8 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	M	W	
	January		
1		**14** Introductions, Administrivia Agents and environments ***Proj 0 out***	
2	[Martin Luther King, Jr. Day **no class**]	**21** Formally specifying problems and problem spaces Informed search ***Proj 0 checkpoint***	*
3	**26** Optimal search Branch and bound A* revisited Admissibility and informedness	**28** Representing game states, moves Minimax revisited	
	February		
4	**2** Evaluating and improving heuristics Real time/time-constrained AI Stochastic, partially observable games	**4** Game day ***Proj 0 due*** ***Proj 1 out***	
5	**9** Alpha-beta pruning Fuzzy logic	**11** Fuzzy logic cont'd ***Proj 1 checkpoint***	
6	**16** Project 1 implementation design Bayesian reasoning revisited	**18** SIGCSE xxxxxx ***Proj 0 followup***	
7	**23** Decision trees, Linear models Entropy and ID3	**25** ID3 cont'd Propositional logic ***Proj 1 due***	
	March		
8	**2** Theorem proving Math proofs Predicate logic Unification ***Exam 1 out***	**4** ***Exam 1*** ***Proj 1 followup***	

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

Wk	M	W
	March	
	SPRING BREAK	
9	**16** Hidden Markov models Noisy channel model ***Proj 2 out***	**18** (cont'd)
10	**23** Supervised learning Training and testing Evaluating models ***Proj 2 checkpoint***	**25** Genetic algorithms
		April
11	**30** Genetic algorithms cont'd	**1** * Neural networks Perceptron models
12	**6** Multilayer neural networks Backpropagation ***Proj 2 due*** ***Proj 3 out***	**8** Ethics of AI classifiers, self-driving cars
13	**13** Genetic-neural programming Planning problems ***Proj 3 checkpoint*** ***Proj 2 followup***	**15** Planning as search
14	**20** Planning algorithms Production systems Intelligent agents	[Research Day **no class**]
15	**27** Responsive agents, emergent systems ***Proj 3 due***	**29** Planning: monolithic systems vs emergent behaviour ***Exam 2 out***
	May	
	Exam 2, Proj 3 followup **Fri, 8 May 8–10:30am**	

* **1 April**: Deadline to withdraw from classes or declare P/F (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

Projects will go out roughly every three weeks. Each will have two checkpoints and a followup that I'll describe in detail elsewhere. Design work is due at the start of class on its due date, and prep work and final versions are due at 8pm on their respective due dates unless I say otherwise. Followup assessments will be in class about a week after the code submission duedate.

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, 7, 4, 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 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 interactions on the Slack channel count. — 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, 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.