

CPSC 415 — Artificial Intelligence

Professor: Stephen Davies
Fall semester 2023

TR 12:30–1:45pm, Farmer 025 (the room formerly known as B7)

Final exam: Thursday, Dec. 14th, High Noon

Office Hours (in Farmer 044):

Mon 11am–2pm
Thu 2–4pm

<http://stephendavies.org/cpsc415>

Artificial Intelligence. . . there’s probably not another term in our field that conjures up as much mystery, excitement, and controversy as this one does. Even outsiders — who would be bored stiff by words like “code,” “transistor,” or “API” — are either thrilled or terrified when the topic of thinking machines comes up. Yet how much is hype, and how much is real? How much potential does AI really have, and what problems yet unsolved can it be used to conquer? How much of a danger really is there in developing such technology? And what does it all mean?

This semester, we’ll embark upon an adventure to find out. Some of what you’ll discover might seem underwhelming, but really shouldn’t be a surprise: any computer system is built of a zillion tiny, interacting parts, all of which have to be painstakingly assembled in just the right way to get anything done. This is as true for a system that claims to be “intelligent” as it is for any other. And what at first seems like “thinking” isn’t really any more than tons of little “dumb” pieces working in concert, when you open up the hood.

That being said, the techniques we’ll learn in this class can be made to do a whoooooole lot, more than perhaps you’ve dreamed of. Whether we’re actually imitating the human thought process, or using an alternate route to get a program to behave *superrationally*, the programmatic structures we’ll create are complex and abstract enough that when they’re all running together, it sort of *does* seem like thinking. In many cases, they can outperform humans at their tasks. It can give you chills.

Roll up your programming sleeves and get ready to plunge into some tough material. Our journey this semester will give you a launching-off point to build the machines, “thinking” or not, of tomorrow.

Goals and Objectives

- To envisage a broad overview of the diverse field of Artificial Intelligence and to discern where various concepts and techniques are situated within it.
- To study some of the important deductive/logic-based AI approaches in depth and understand how they work under the hood.
- To design and implement AI agents that can accomplish simple goals and illustrate certain important techniques in action.
- To confront some of the philosophical and ethical issues surrounding the enterprise of Artificial Intelligence, and figure out where you stand on the key controversies.

Student Learning Outcomes

Upon completing this course, students will be able to:

- Articulate the scope of the diverse field of Artificial Intelligence and where various concepts and techniques are situated within it.
- Implement important deductive/logic-based AI approaches and explain the trade-offs inherent in their implementation.
- Design AI agents that can accomplish simple goals and attempt to optimize their scores on quantitative performance measures.
- Characterize the important philosophical and ethical issues surrounding the enterprise of AI, and to defend positions on the key controversies.

Rules of the game

1. There are NO stupid questions. I will never belittle you or make you feel dumb about anything. Your job is not to already know everything before you start, but to roll up your sleeves and work hard to try and learn. Even if your question is, “Stephen, I totally didn’t even get that, can you start over and explain it again?”, please ask.
2. This class will be interactive. When I point at you in class, say your first name, and be prepared to try and answer questions. (Don’t worry if you don’t know “the answer” every time. Go ahead and give it a shot, or else just say “**pass**” which is always accepted.)

3. Exams and quizzes will be based solely on the material I cover in class. However, the suggested book readings will reinforce these concepts for you in tremendous ways, so I suggest you make sure to learn both!
4. Don't skip class. Just don't. It's bad form. I work hard on preparing for class, to make it compelling and relevant. It hurts my feelings when you don't come. Plus you miss out on important stuff, and you'll end up falling behind if you skip lecture. So come every time. Come happy, fresh, excited, ready to think and to participate.
5. **Please, no laptops during class.** I've had students claim that they take notes on their laptop during lecture, but even if it's true, those things are too big a distraction to you and your fellow students to make it worth it. Just stay tuned in, because I move fast.

The Honor Code and this course

For this course, **all the work that you turn in for a grade must be solely your own work, period.** Specifically, this means:

- Quizzes and exams must be taken alone, in a quiet place, without any form of contact with anyone. Unless otherwise stated, you may not consult any written source or Internet site to complete them.
- **You must write all your own programs in their entirety.** I don't mind if you chat informally about the programming assignments with your fellow students, but you must not show anyone else your code nor look at anyone else's code. This includes people who are not in the class nor even at UMW. **I am happy to help you over email or in office hours about whatever questions and problems you have.**
- For the responsive readings, you must write all your own text. You may not copy it from anywhere or get anyone to write any part of it for you. However, you *are* permitted (and encouraged) to go to the UMW Writing Center for assistance on your writing. Also, you are **only** allowed to use an AI assistant (like ChatGPT) **in the specific way I will describe in class.** (It will involve a documented draft-and-revision cycle, all parts of which you will turn in.)

Note that the Russell and Norvig book is world-famous and widely used. Hence, there are all kinds of solutions to all kinds of problems available on the Internet. **It is an Honor Code violation to look for or consult any of these.**

Movie nights

We'll be having periodic movie nights this semester in the HCC Digital Auditorium. These will involve viewing a film featuring some aspect of artificial intelligence, eating popcorn and other snacks, and participating in an after-film guided discussion. You will also have a chance to earn additional XP at these movie nights: +3XP just for showing up, and +1XP for making at least one meaningful comment during the after-viewing discussion.

Note that these films are being showed for educational purposes only, in order to form the basis for an academic discussion afterwards. Also: attendance is restricted to only students enrolled in the course.

I'll be scheduling these film viewings on a variety of different weeknights, so as to accommodate people's schedules. They will start promptly at **7pm**. You may have a time conflict for some of them; hopefully not for all of them:

Thu Sep 7	<i>Ex Machina</i> , Alex Garland (2014)
Wed Sep 13	<i>Her</i> , Spike Jonze (2013)
Mon Sep 18	<i>Battlestar Galactica</i> , season 1 (excerpts) (2005)
Tue Oct 3	<i>Blade Runner</i> , Ridley Scott (1982)
Tue Oct 10	<i>Blade Runner 2049</i> , Denis Villeneuve (2017)
Mon Oct 23	<i>Star Trek</i> : (selected episodes) (1968, 1998)
Wed Nov 1	<i>2001: A Space Odyssey</i> , Stanley Kubrik (1968)
Tue Nov 7	<i>Westworld</i> , season 1: "pilot," Jonathan Nolan (2016)
Thu Nov 16	<i>Westworld</i> , season 1: "The Stray," Neil Marshall (2016)
Wed Nov 29	<i>Terminator 2: Judgment Day</i> , James Cameron (1991)
Mon Dec 4	<i>The Matrix</i> , Lana and Lily Wachowski (1999)

All movie showings in HCC Digital Auditorium.

Book (recommended)

Russell and Norvig, *Artificial Intelligence: A Modern Approach*, Fourth Edition, Pearson, 2022.

This is hands-down the world's greatest book on AI, and one of the best Computer Science books ever written, period. It is broad, deep, well-anchored in the literature, and perhaps most importantly: fantastically well-written. The authors are two of

the top names in the field, and they have invested enormous effort into making this one of the most fascinating texts you'll ever read. It's well worth every moment you spend on it, this semester and beyond.

As I'll explain in class, AI can be thought of as encompassing two big subfields: the deductive/logic-centric half, and the inductive/data-centric half. It's really much too big for a single course. The good news, however, is we already have a course at UMW (CPSC 419) which focuses on the second of these two halves (essentially, the material in chapters 18 to the end of the book). Therefore, we're going to focus on the first half of the field and of the book this semester: chapters 1 through 17.

Submitting programs and things

Some of your assignments this semester will be submitted electronically, and some by hardcopy. For the former, I'll give special instructions for each assignment. Most will involve sending me an email with an attachment that is either a Python file or a git bundle with your project repo. **I do insist that you *exactly* perform the submission procedures I specify.** This includes adhering to the specific subject line I specify so I can readily distinguish submissions from my hordes of other email.

For hardcopy assignments, stick them in my "415" manilla folder hanging outside my office door.

Earning XP

There will be opportunities to earn XP throughout the course. Some of these will be spontaneous as the mood strikes me. Some you can earn by completing in-class activities. Some may be in response to impressive things I see you do as the semester progresses. The following opportunities are *guaranteed* to be available to you:

Guaranteed XP opportunities:

Activity	Possible XP
Eight quizzes	30 each
Six programming assignments	50 each
Four responsive readings	20 each
Final exam (comprehensive)	50
Various and sundry others	varies

Grading

Grading this semester will be based on “experience points” (XP). As you complete activities, you will earn XP towards your final total. XP can never be lost, only gained, but you have to earn what you get (*i.e.*, you don’t “start off with a 100%” and lose points based on mistakes you make).

Here are the levels you may achieve, together with their semester grade and the XP requirements for each:

Level	Total XP	Semester grade
Agent Smith	650	A+
Ava	600	A
Ash	560	A-
BB-8	530	B+
Bishop	500	B
Baymax	470	B-
C-3PO	440	C+
Cortana	420	C
Cylon	400	C-
R2-D2	380	D+
Data	360	D
Dolores	340	D-
Terminator T-800	320	.
HAL-9000	300	.
Ultron	280	.
Maeve	260	.
K-2SO	240	.
#6	220	.
TARS	200	.
M3gan	175	.
Rachael	150	.
L3-37	125	.
Eve	100	.
Wall-E	75	.
Rosie the Maid	50	.
Power droid	20	.
Dead battery	0	.

Basis for determining mid-semester reports

If you get a “U” (unsatisfactory) for this class on your mid-semester report, it means that I’m giving you a red flag that you’re underachieving in the class. Normally that means you’re not engaged, and that you have turned in late (or not at all) multiple assignments. If you don’t get a “U”, that doesn’t necessarily mean you’re ace-ing the class: it just means I chose not to sound a big alarm.

Late policy

No late work will be accepted this semester. Get your stuff in on time: there’s no excuse not to!

Guidelines for class participation

I believe that students learn best when they participate wholeheartedly in all aspects of the learning process. Hence while your grade will not be determined by any “class participation score” *per se*, it is very much to your advantage, and very much recommended, that you join in during class discussions, ask questions, and make comments.

Btw, I do sometimes give out XP cards to students for particularly well-framed questions or analyses in lecture.

How to reach me

Come to office hours, see me after class, or e-mail me (stephen@umw.edu)!

How to reach you

I will post announcements to the course website often, so be sure to subscribe to its RSS feed or check it manually at least once a day! Also, I will occasionally be communicating with you outside of class time via e-mail, so make sure to check your UMW e-mail every day!

Title IX Statement

UMW faculty are committed to supporting students and upholding the University's *Policy on Sexual and Gender Based Harassment and Other Forms of Interpersonal Violence*. Under Title IX and this Policy, discrimination based upon sex or gender is prohibited. If you experience an incident of sex or gender based discrimination, we encourage you to report it. **While you may talk to me, understand that as a "Responsible Employee" of the University, I must report to UMW's Title IX Coordinator what you share.** If you wish to speak to someone confidentially, please contact the confidential resources below. They can connect you with support services and help you explore your options. You may also seek assistance from UMW's Title IX Coordinator; their contact information can be found below. Please visit <http://diversity.umw.edu/title-ix/> to view *UMW's Policy on Sexual and Gender Based Harassment and Other Forms of Interpersonal Violence* and to find further information on support and resources.

Resources

Ruth Davison, Ph.D.
Title IX Coordinator
Lee Hall, Room 401
540-654-5656
rdavison@umw.edu

Confidential Resources

On-Campus

Talley Center for Counseling Services
Lee Hall 106, 540-654-1053

Student Health Center
Lee Hall 112, 540-654-1040

Off-Campus

Empowerhouse (24-hr hotline)
540-373-9373

RCASA (24-hr hotline)
540-371-1666

Recording Policy

Classroom activities in this course may be recorded by students enrolled in the course for the personal, educational use of that student or for all students presently enrolled in the class only, and may not be further copied, distributed, published or otherwise used for any other purpose without the express written consent of the course instructor. All students are advised that classroom activities may be taped by students for this purpose. Distribution or sale of class recordings is prohibited without the written permission of the instructor and other students who are recorded. Distribution without permission is a violation of copyright law. This policy is consistent with UMW's *Policy on Recording Class and Distribution of Course Materials*.

Accessibility statement

The Office of Disability Resources has been designated by the university as the primary office to guide, counsel, and assist students with disabilities. If you receive services through the Office of Disability Resources and require accommodations for this class, please provide me a copy of your accommodation letter via email or during a meeting. I encourage you to follow-up with me about your accommodations and needs within this class. I will hold any information you share with me in the strictest confidence unless you give me permission to do otherwise.

If you have not made contact with the Office of Disability Resources and have reasonable accommodation needs, their office is located in Seacobeck 005, phone number is (540) 654-1266 and email is odr@umw.edu. The office will require appropriate documentation of disability.

Basic needs security

Learning effectively and engaging wholly in class is dependent upon our basic security and having our fundamental needs met: having a safe place to sleep at night, regular access to nutritious food, and some assurance of safety. If you have difficulty affording groceries or accessing sufficient food to eat every day, or if you lack a safe and stable place to live, please contact Chris Porter, Assistant Dean of Students, at cjporter@umw.edu. Additionally, the Gwen Hale Resource Center is a free resource on campus, providing food, toiletries and clothing to any member of our community. It is open Monday, Tuesday and Friday from 1pm-6pm, on the 5th floor (floor A for Attic) of Lee Hall, or resource@umw.edu. Finally, you are always welcome to talk with me about needs, if you are comfortable doing so. This will enable me to provide any resources I may possess.

Road map

The calendar for the course, complete with assignment due dates, tests, etc., will be maintained on the course website at <http://stephendavies.org/cpsc415>. Visit there often. It is guaranteed to change!

As of the first day of class, however, here's a rough preview of my best guess at how things will go:

Week	Topics	Chapters	Due
1	Background; the AI architecture	1–2	P0: Python
2	Problem-solving agents	3	P1: Agents & environments
3	Advanced search techniques	4	RR1: Foundations
4	Advanced search techniques	5	P2: Intelligent search
5	Constraint satisfaction problems	6	RR2: Applications
6	Agents based on Propositional Logic	7	
7	Agents based on Predicate Logic	8	P3: Adversarial search
8	Agents based on Predicate Logic	9	RR3: Societal impact
9	Planning agents	10	
10	Planning; knowledge representation	11–12	P4: Logic agents
11	Mathematical probability	A.3, 13	
12	Probabilistic agents	14	P5: Planning agents
13	Probabilistic agents	15	RR4: Ethics
14	Decision-making agents	16	
15	Decision-making agents	17	P6: Probabilistic agents

“P”=Program, “RR”=Responsive Reading