



University
of Vermont

College of Engineering and
Mathematical Sciences
Computer Science

CS 2510/ Introduction to Artificial Intelligence / 2026 Fall

Compiled: 2026-09-03 07:15

Instructor: Clayton Cafiero

Email: cbcafier@uvm.edu

Office: Innovation E309

Office hours:

M 02:30–03:30 PM, Th 09:30–10:30 AM,

F 09:30–10:30 AM, or by appointment

Website: <https://www.uvm.edu/~cbcafier>

GTA: TBA

Email: TBA

Office hours: TBA

Location: TBA

UTA: Layla Musallam

Office hours: TBA

Location: TBA

B: M/W/F 08:30–09:20 AM, Lafayette L411

Welcome to the course!

Overview and learning objectives

Introduces the foundations of artificial intelligence, its history and selected methods. Topics include: logic- and rule-based approaches, knowledge representation, heuristic search, A*, IDA*, α - β pruning, minimax, expectiminimax, Markov models and Markov decision processes, decision tree, ensemble learning / random forest, the neural model and simple multi-layer perceptrons, and brief introduction to modern AI. Other topics, if any, TBD.

Prerequisites

CS2240 with grade of C- or better.

Learning objectives

Through lecture, readings, discussion, in-class active learning exercises, and homework, students will be able to:

- describe in general terms the history and development of the field of artificial intelligence;
- use Prolog to develop small scale knowledge-bases of facts and rules, and then use queries and helper functions to produce inferences / predictions (including application of recursive rules);
- construct and evaluate simple knowledge representations, and be able to identify how a knowledge representation determines what inferences are sanctioned;
- identify common biases in knowledge representations;
- understand how heuristic search can outperform naive search in many instances;
- implement and use A* and IDA* in code;
- choose or design an appropriate heuristic for use with A* and IDA* algorithms (also ensuring the heuristic is admissible, and if possible, consistent);
- apply α - β pruning to heuristic search, understand the benefits of doing so, and, given small problem instances, identify subtrees which can be pruned;
- compute best moves / actions, using minimax or expectiminimax, given a game tree for evaluation;
- compute a stochastic matrix (Markov model) from a set of observations and use the resulting model to make single- and multi-time step predictions;
- understand the basics of Markov decision processes (MDPs) and how policies are devised to maximize gain, given a particular environment with rewards and actions;
- apply the Bellman equation in the context of MDPs;
- construct decision trees using conventional algorithms and produce classification predictions using such trees;
- understand how entropy and information gain are used in constructing decision trees, and perform such computations by hand or in code;
- implement random forest in code and compare results with individual decision trees;
- understand and compute metrics commonly used to evaluate model performance (*e.g.*, precision, recall, accuracy, F-measure, *etc.*);
- understand how simple neurons (units) aggregate inputs and produce output via various activation functions, and perform simple computations on small problem instances (for example, applying the perceptron learning rule);
- perform experiments to test how hidden layers, non-linear activation functions, and learning rate affect model performance and training time; and
- describe the concepts behind backpropagation without detailed or rigorous mathematical treatment.

Course materials

Materials including readings, lecture notes, and papers of historic significance will be circulated via Brightspace.

Software

Students should have a working installation of Python (Python 3.12 or higher is recommended). We will use Scikit-Learn, TensorFlow 2 or PyTorch, and Matplotlib. All of these are pip-installable (instructions will be circulated via Brightspace). With these installed Numpy, Scipy, and many other packages will be installed as dependencies. A new virtual environment is *strongly* recommend to manage and isolate requirements.

We will do some programming in Prolog. Installation will vary from machine to machine, from OS to OS, and there is no universal recipe. See SWI-Prolog: <https://www.swi-prolog.org/>. One's best bet on systems with package managers (Linux, macOS) is to use the package manager (*e.g.*, Apt, Aptitude, dpkg, Homebrew, MacPorts, *etc.*) If installation is problematic, students can run Prolog programs in the SWISH web interface at: <https://swish.swi-prolog.org/>.

Correspondence

Please use email for electronic correspondence (and not MS Teams). As I teach multiple courses please indicate the course in which you are enrolled in the subject line. Please use your UVM email for all correspondence.

Weekly Schedule of Topics (Tentative and Subject to Change):

Week	Topic
1	Introduction; history of AI; logic foundations; Prolog
2	Logic programming with Prolog
3	Knowledge representations and ontologies; expert systems
4	Heuristic search, A*; admissibility and consistency
5	Heuristic search, IDDFS and IDA*
6	Adversarial games, minimax, and α - β pruning
7	Midterm exam
8	Non-deterministic games and expectiminimax; quiescence and null move search
9	Markov models and Markov decision processes
10	Decision trees, random forest, and ensemble learning
11	The neural model, McCulloch-Pitts, Hebb, perceptrons
12	Multi-layer perceptron networks (neural networks)
13	AI in its current state: LLMs, agents, <i>etc</i>
14	Additional topic(s) TBD

In addition to the above, we'll read a few papers and set aside some class time for (hopefully lively) discussion.

Grading

Weight	Assessment
10%	in-class discussion and active learning
45%	homework and mini-projects
5%	scribing
10%	quizzes (11, drop one)
10%	midterm exam
20%	final exam
100%	TOTAL

Grades will be assigned using a conventional scale:

≥ 98*	A+	≥ 77	C+
≥ 93	A	≥ 73	C
≥ 90	A-	≥ 70	C-
≥ 87	B+	≥ 65	D
≥ 83	B	< 65**	F
≥ 80	B-		

* To earn an A+ you must reach or exceed the indicated threshold excluding bonus points.

** To pass the course you must reach or exceed the indicated threshold excluding bonus points.

Scores will not be rounded up when assigning letter grades. Moreover, I reserve the right to reduce your final score by one full percentage point if you ask to have your score rounded up.

Exams

There will be a midterm examination administered in class (before fall recess).

The Office of the Registrar schedules final exam dates.

- M 2026-12-14 07:30 AM–10:15 AM, Lafayette L411

See: <https://www.uvm.edu/registrar/final-exams> for complete schedule.

Scribing

In most weeks a group of two or three students, working independently, are expected to take notes in class and write them up in electronic form for distribution to the rest of the students in the course. Since I'm sure you're all taking notes anyway, I don't think you'll find this too onerous. Notes are

to be prepared in Markdown format, with inline LaTeX as needed for equations and mathematical notation. You should be present when scribing, since the objective is to capture what was actually transacted in class. Scribing after the fact must necessarily rely on sources other than in-class lecture, so this is not permitted. Unexcused absence during your week of scribing will necessarily result in a lowered maximum score for your scribing submission. Instructions and guidelines will be posted on Brightspace, and notes will be submitted via Brightspace.

I will assign groups at random, and you should inform me as soon as possible if you have a conflict and we'll see what we can do to swap assignments for a given week. Scribing will commence in week two, and will not take place in weeks in which we have an in-class exam. Look on Brightspace for your scribing assignment.

If you have a documented accommodation which entitles you to a peer note taker (through the ACCESS office), you are exempt from scribing, but I will give you a small supplemental assignment you can do on your own instead. In the event we have a need for a peer note taker (for the ACCESS office), anyone who volunteers and serves as a peer note taker will be exempt from other scribing duties and will receive full credit.

Academic integrity

The Department of Computer Science enforces UVM's Code of Academic Integrity. Any suspected violation of this policy will be referred immediately to UVM's Center for Student Conduct (<https://www.uvm.edu/sconduct>). Sanctions for a violation may include a grade of XF in the course. Additional violations can result in dismissal from the university. In a word: *Don't*. All students should read and understand this policy. See: <https://go.uvm.edu/cai>.

Collaboration on quizzes and exams is strictly prohibited. Use of online services as a source of solutions is strictly prohibited. Using generative AI such as ChatGPT or Claude, or websites such as Chegg or Course Hero to complete coursework is a form of academic dishonesty. Work you submit for an individual grade must be your own. Any work not produced by you must be cited. For certain assignments, students may collaborate on homework (typically limited to teams of two). If you collaborate with another student on an assignment, be sure to indicate team members as specified. If you have any questions ask!

Any attempt to tamper with or defeat any autograder is a form of academic dishonesty. This applies wherever autograders are in use, for example on Brightspace or Gradescope.

All code submitted by students is subject to code similarity review.

Exams, quizzes, homework assignments, answer keys and solutions, presentations or lecture notes, specifications and rubrics are copyright protected works, unless clearly and explicitly indicated otherwise. Any unauthorized copying or distribution of protected works is a violation of federal law and may result in disciplinary action. This includes submission of protected works as prompts to generative AI. Sharing of course materials without the specific, express approval of the instructor may be a violation of the University's Code of Academic Integrity and an act of academic dishonesty, which could result in disciplinary action. Violations will be handled under UVM's Intellectual Property Policy and Code of Academic Integrity, as appropriate. See: <https://go.uvm.edu/ipp> and <https://go.uvm.edu/cai>.

Late policy / extensions

Each homework assignment has a specific due date / time. You may submit work up to 24 hours after the due date / time, however, late submissions will be penalized 20%. Submissions that are more than 24 hours late will not be accepted unless an extension has been granted. We will consider reasonable requests for extensions when extenuating circumstances arise. (It can't hurt to ask.) However, extensions will not be granted if the request for extension is made within 24 hours of the time an assignment is due, except in the most extraordinary circumstances. So if you wish to request an extension, *do so early!* If an extension is granted, you must submit your work by the agreed-upon extension date.

Student course evaluations

Students are warmly encouraged to complete an evaluation of the course at its conclusion. Evaluations are anonymous and confidential, and the information gained, including constructive criticisms, will be used to improve the course.

The secret word is “sympathetic.”

Accommodations

In keeping with UVM policy, if you have a documented disability and are interested in utilizing ADA accommodations, you should contact Student Accessibility Services (SAS), the office of Disability Services on campus for students. SAS works with students and faculty in an interactive process to explore reasonable and appropriate accommodations, which are communicated to faculty in an accommodation letter.

Contact SAS: A170 Living/Learning Center; +1 802 656 7753; access@uvm.edu; or visit <https://www.uvm.edu/access>.

If you are entitled to use the Exam Proctoring Center, please book reservations at least four days in advance.

Class participation

You are expected to be an active participant in class. The more engaged you are, the more you will learn—and the more fun you'll have. This includes being prepared and attentive, responding when called on, participating in group discussion, and asking questions as appropriate. When it comes to asking questions, *please don't be shy!* There's no such thing as a “dumb” or “silly” question. If there's something you don't understand—*ask!* Asking questions helps you understand the material presented in the course. Asking questions is good for your classmates. It's almost certain that if you need clarification on some point, that there's at least one other student in the class with the same question. So help each other out—*ask!* Finally, when you ask a question you help the instructor to do a better job of explaining. If someone explains something, and you still don't quite grasp it, it's not unlikely that the explanation could be improved or clarified.

Inclusivity and bias

UVM is a place where you should be treated with respect and kindness. We welcome individuals of all ages, backgrounds, beliefs, interests, ethnicities, genders, gender identities, gender expressions, national origins, religious affiliations, sexual orientations, ability, and other visible and non-visible differences. All students are expected to contribute to a respectful, welcoming and inclusive environment for every other member of the community. If you ever feel that you have been unfairly treated or judged by an instructor, a mentor, another student, or another member of the community, please let someone know. Your instructors and advisors in the CEMS Office of Student Services are available to discuss any concerns, or you can report an incident of bias through the bias report program (<https://go.uvm.edu/brp>).

Conduct

Be kind to one another and to yourself. Be respectful of yourself, others, and the institution. Please arrive on time. Please, no food in class. Please, no cell phones in class (except for medical monitoring with approved accommodations; see: attendance policy for more on this). You may use a laptop or tablet, but only for active learning sessions, pair programming, taking notes, or assistive technologies. Recording or capturing audio or video of class sessions is prohibited. Still photos of whiteboard / chalkboard are OK but discouraged. Please ask first so I can get out of frame.

For other policies on classroom conduct, please see: <https://go.uvm.edu/srr> and <https://go.uvm.edu/csc>.

Promoting health and safety

If you are concerned about a UVM community member or are concerned about a specific event, we encourage you to contact the Dean of Students Office (+1 802 656 3380). If you would like to remain anonymous, you can report your concerns online by visiting the Dean of Students website at <https://www.uvm.edu/deanofstudents>.

Wellbeing resources

- Center for Health and Wellbeing: <https://www.uvm.edu/health/services>
- Counseling and Psychiatry Services (CAPS): +1 802 656 3340
- Food Insecurity Assistance: <https://www.uvm.edu/health/food-insecurity-uvm>

Student advocacy

https://www.uvm.edu/deanofstudents/student_advocacy

Your identity at UVM

Students at UVM can specify the first name and pronoun they want used on campus. For information on how to update your preferred name and personal pronouns as well as keeping your legal name private, and UVM policy on lived name and pronouns, see: <https://go.uvm.edu/lmpr>.

Religious holidays

Students have the right to practice the religion of their choice. In order to receive extensions or excused absences, you should submit via email your documented religious holiday schedule for within the first week of the course. Reasonable extensions will be granted where assignment deadlines conflict with religious holidays.

Student athletes

In order to receive extensions or excused absences, you should submit via email appropriate documentation as soon as possible, preferably within the first week of the course. Reasonable extensions will be granted where assignment deadlines conflict with team events or team travel.

Statement on alcohol and other drugs

We want you to get the most you can out of this course. Therefore, you are expected to familiarize yourself and abide by the University's policies with regard to alcohol, cannabis, tobacco, and other drug use. See: <https://go.uvm.edu/actd>. Please do everything you can to optimize your learning and to participate fully in this course.

Class format changes

The University of Vermont reserves the right to make changes in the course offerings, mode of delivery, degree requirements, charges, regulations, and procedures contained herein as educational, financial, and health, safety, and welfare considerations require, or as necessary to be compliant with governmental, accreditation, or public health directives.

Changes to this document

This document is subject to change. Any such change will be communicated via an announcement on Brightspace. The latest version of the syllabus will always be available on Brightspace.

Policy regarding use of AI

Here AI is understood to mean any generative AI, which includes but is not limited to ChatGPT, Claude, Copilot, Llama, DeepSeek, Gemini, Kimi, *etc.*. It also includes AI summaries in web search results, and tools such as Cursor, Windsurf, Replit, Codex, ZenCoder, *etc.*, and personal intelligence assistants such as Siri / Apple Intelligence, Alexa, Perplexity, and others.

1. Exams and in-class assessments. Use of AI is *strictly prohibited* as a means of researching, producing, or verifying answers in exams, quizzes, essays, and other in-class assessments. Any use of AI in such assessments will result in immediate referral to the Center for Student Conduct. *No exceptions.*

Why? Exams and in-class assessments are intended to assess your knowledge and ability to apply concepts you've learned in class. These are not intended to assess your ability to use AI to produce answers. Moreover, should one student use AI and not another, the student using AI might gain an unfair advantage over others.

2. In-class discussions. Use of AI is prohibited for use in in-class discussions, unless explicitly permitted by the instructor at a specific time. If you are found to be using unsanctioned AI in an in-class discussion, you will be asked to stop, and you'll receive a zero for participation if the discussion is included in grading. A second infraction will result in referral to the Center for Student Conduct.

Why? In-class discussions are opportunities to ask questions and share your ideas with classmates and the instructor. We're interested in what *you* think, not what a chatbot says. Moreover, studies have shown that use of AI in such contexts can result in a reduced diversity of opinions expressed. Some participation is included in grading, and so, again use of AI can give a student an unfair advantage over others.

3. Homework. Use of AI is *strictly prohibited except where noted (below)* as a means of researching, producing, or verifying answers or responses in homework *except where homework specifically calls for the use of AI*. As you are no doubt aware, this is difficult to enforce. Nevertheless, you should not use AI for assistance in homework. *You should be able to read, understand, and explain in detail any other work product you may submit for grading. Do not submit anything you don't completely understand.* Moreover, should one student use AI and not another, the student using AI might gain an unfair advantage over others. If I have questions about how your work was produced I reserve the right to have you come to a meeting in which you give an account of your work and how it was produced. In such a meeting, you should be able to answer questions without hesitation. If I determine that AI was used to produce homework, in whole or in part, I will refer the case to the Center for Student Conduct.

Why? Homework is an opportunity for you to engage deeply with course materials and concepts, and it is meant to provide you with challenges. It is through these challenges that we learn and cultivate our reasoning and problem solving skills. Without human effort, you reduce homework to producing a mere artifact. That is not at all the objective—education is not about producing artifacts, it's about *learning*. You will retain more if you fully engage in problem solving and analysis. What's important is what goes into producing the work—What do you know? How can you apply what you know? How can you strengthen your problem solving skills?

What are the exceptions? Some homework and projects may require the use of AI. Such cases will be clearly indicated, and guidelines for permitted and prohibited use will be supplied in the assignment specification and rubric.

While the use of generative AI for *producing* prose—reflections, essays, or similar writing—is prohibited, it is permissible to use AI as a proofreader for writing you have produced. UVM has a site license for Copilot, which is integrated into UVM-supplied instances of Microsoft Word. This is very much preferred to outside, third-party tools. If you do use something other than UVM’s Copilot instance you must make a note of it in a comment, footnote, or end note. Even if you use Copilot or other tool, you should *never* use it to produce a first draft. Preserve your voice—make sure your words are indeed your own.

If you have a documented disability (e.g., dysgraphia) or English is not your native tongue and this presents significant difficulty in writing, please contact me about an accommodation.

4. Autocomplete in application software and IDEs. What about the role of autocomplete? I recommend turning it off.

Why? In many contexts, *autocomplete defeats its own purpose*. Rather than assisting, it can add to the cognitive load required to complete a task. Rather than just writing the sentence or line of code we had in mind, we may be faced with the challenge of evaluating autocomplete suggestions—which can be a great way to get sidetracked. Some autocomplete features offer multiple suggestions which only compounds the difficulty.

Furthermore, recent studies have shown that in writing autocomplete can subtly steer our thinking and effectively put words in our mouths. This is of concern in light of the fact that the companies which develop and deploy these tools have agendas of their own, and autocomplete can be an insidious way to alter people’s views.

That said, this is only a recommendation. Nevertheless, please pay attention to any cognitive burden that autocomplete may cause, and do everything you can to preserve your own voice in writing.

5. AI as an explainer of things. It may be the case that you find certain course content confusing or challenging, and there are some legitimate ways one might use AI to explain concepts, provide additional problem instances for practice or limited coaching on various topics. We refer to the beneficial aspects of this use as *cognitive augmentation*. It’s OK to use AI in this way, so long as it is of real benefit the learner (you).

However, there’s a problem here. If students use AI in this way, it circumvents the usual avenues for getting help—office hours with an instructor or a teaching assistant. By going directly to AI for help, instructors lose valuable feedback. For instance, if many students ask for help on a specific topic then the instructor is made aware that some students need help. This can provoke questions and actions on the part of the instructor: Maybe I need to explain X (whatever X is) in a different way. Maybe we need to add some think-pair-share exercises or other check-in in class. Maybe I need to spend an extra day on topic X. By going straight to AI, the challenges students face become invisible to the instructor (until exam time, and then it’s too late). Another problem is that AI advice may not be aligned with actual course content and learning objectives. But unlike office hours, AI is available 24/7/365.

Accordingly, here’s the policy on such assistance. I will make an *anonymized open drop box* available for information students wish to share in this context. This could include information on topics you used AI to explain, actual prompts, or even entire transcripts of chat sessions. These will be anonymous and ungraded, but it might help bridge the gap between student’s need for additional explanation and the need for the instructor to receive actionable feedback.

While not AI, the use of YouTube, Computerphile, 3Blue1Brown, MIT Open Courseware, or other platforms for delivery of instructional material also disrupts the feedback loop that instructors need to continuously improve instruction. Therefore, I encourage you to submit links to YouTube, Computerphile, 3Blue1Brown, MIT Open Courseware, or other resources you have found helpful to the anonymized drop box.

6. **Note-taking.** Use of AI for note-taking is prohibited. This includes Otter.ai and other AI transcription and note-taking tools.

Why? Three reasons. First, if you rely on note-taking and transcription tools, you won't remain as attentive in class, and you'll likely remember less than if you had been engaged and took your own notes. Second, these tools consume streaming data and I do not consent to have class recorded or streamed without my expressed consent. Third, if other students are aware that AI note-taking is being used, they may feel uncomfortable asking questions, offering answers to my questions, or expressing their opinion. This would result in less engaging class sessions for all.

7. **Copyrighted material.** Course materials and assignments are copyrighted unless otherwise noted. Even open source materials can be copyrighted. As a matter of policy, please don't copy/paste or upload copyrighted materials to AI tools.

Why? Uploading copyrighted material to AI tools risks legal infringement, violation of terms of service, and loss of intellectual property rights, as many platforms store and use uploaded data to train their models. This practice can lead to proprietary content being reproduced by the AI for other users, constituting potential copyright violation or even violation of an open source license which requires attribution.

8. **Privacy and FERPA.** As we will have group exercises and may have group homework, it's important that you never share identifying information with AI. If you do use AI for proofreading, make sure all names are expurgated before sharing any content with AI tools.

Why? Sharing identifying information about any student other than yourself may constitute a violation of FERPA (Family Educational Rights and Privacy Act).

9. **General.** Any unauthorized use of generative AI, or uncited use where permitted, constitutes a violation of UVM's Code of Academic Integrity and Code of Student Conduct. Violations or suspected violations will be referred to the Office of Student Conduct.

Why? This is a matter of UVM policy. See UVM's Code of Academic Integrity and Code of Student Conduct for more.

10. **Instructor responsibilities.** In keeping with the expectations for students, I will be sure to follow best practices in my own work and in any course materials including lecture notes, handouts, worksheets, exams, quizzes, *etc.* I will indicate clearly where and when AI has been used (if at all) in revising, proofreading, fact-checking, or other contexts. I guarantee that I will never use AI to provide feedback to you on any assessments, ever.

Why? This is necessary for transparency, integrity, and pedagogical benefit. Students have a right to know if AI is used in developing any teaching materials. Failure of the instructor to adhere to the same principles that guide these policies would be hypocritical and dishonest.

11. **Right to opt out of AI use.** There are many ethical issues surrounding the use of generative AI. If you are uncomfortable about using generative AI, your views are welcome and respected, and you are in no way obliged to use it.

Explicitly excluded from the above are the models we build or use in class or homework. Engagement with these is essential to achieving course learning objectives and acquiring AI literacy skills.

12. **Questions and comments.** If you have questions about generative AI or this policy, please do not hesitate to ask.

WCAG 2.1 AA-level standards

This document conforms to WCAG 2.1 AA-level accessibility standards as required by UVM policy.

AI usage statement

No AI was used in researching, planning or writing this document.

However, AI (Anthropic Claude Code 2.1 / Sonnet 4.6) was used to review output of accessibility verification tools (veraPDF (<https://verapdf.org/>)) and to make recommendations to address compliance gaps with respect to WCAG 2.1 AA-level accessibility standards. The specific recommendations provided by Claude were:

- enabling LaTeX document metadata tagging in document preamble,
- switching from XeLaTeX to LuaLaTeX as compiler, and
- replacing nested tabulars with `p{}` columns in tables.