



## CS 580-002 Introduction to Artificial Intelligence — Spring 2023

Tuesday & Thursday 7:20-10:00pm, [Planetary Hall](#) room 129. See <https://cs.gmu.edu/~hrolenok/teaching/cs-580-spr2023/index.html> for current syllabus. Be sure to check your @gmu.edu email and blackboard regularly for updates.

**Instructor:** Brian Hrolenok  
**@gmu.edu email:** hrolenok  
**Office Hours:** TBA

**TA:** TBA  
**@gmu.edu email:**  
**Office Hours:**

### Note about compressed schedule

This course is being offered on a compressed 7.5 week schedule. This means that several important dates may be different if you are familiar with the regular semester at GMU. Specifically, our first meeting will be on Tuesday, March 21st, and our last class will be on Thursday, May 4th. This is session II, see the [registrar's calendar](#) for more details. Additionally, our class meets twice as frequently as compared to the full semester (twice a week, for 180 minutes per meeting). This is a **significant amount of work**, so please be sure you are prepared to commit the time required.

### Course description

CS 580 - Introduction to Artificial Intelligence is a 3-credit introductory course intended for graduate students. Artificial Intelligence is subfield of Computer Science which covers the design, implementation, and analysis of computational systems that can be said to reason, learn, or act rationally. This course presents a broad overview of this material using an agent based approach, and has a particular focus on the details of implementation. The class is programming intensive, and a strong background in the programming language of the assignments (Python) will be very useful.

#### Learning objectives:

- To provide a broad survey of algorithms and approaches in AI.
- To develop a deeper understanding of several major algorithms in AI.
- To develop the design and programming skills that will help you to build computational artifacts that act intelligently.
- To develop the basic skills necessary to pursue research in AI.

**Prerequisites.** The recommended prerequisites for this course are CS 310 and CS 330, and familiarity in the following topics and strong mathematical maturity will be useful:

- Probability
- Linear algebra
- Calculus

**Textbook:** *Artificial Intelligence: A Modern Approach* by Stuart Russell and Peter Norvig.

The readings given below are for the third edition, section numbers may be different for other editions. Ideally, readings should be completed before class so you can bring your questions to lecture.

### Homework and Problem Sets

This semester, there will be two different kinds of take-home assignments: *homework* and *problem sets*. **Homework** assignments are programming focused, and come with an autograder for you to test your code. They are graded for correctness with the autograder and by manual inspection. Their purpose is to give you hands on practice implementing the algorithms and techniques discussed in class in a concrete setting, and, specifically, experience with the unique challenges of debugging in these domains.

**Problem sets** are more focused on understanding and exploring the lecture/reading topics from a theoretical and mathematical standpoint. They involve short-answer responses, proofs, open-ended exercises, and occasionally some short programming. These assignments **are not** graded for correctness and they have **no autograder** — any submission before the due date will receive full credit. The purpose of these assignments is to keep you synchronized with the course lectures and readings, give you touchstones for potentially challenging concepts that do not easily fit with the other programming based homework, and to expand on topics that we may not have time to cover in depth during lecture. Some solutions or solution approaches may be covered in lecture, time permitting. Problem set submissions should be a single PDF file unless specified otherwise in the assignment description.

All assignment submissions will be handled through [Blackboard](#), and are due by the date and time listed there. Submissions by email will not be accepted.

### Collaboration

All of the assignments in this course can be completed individually or in groups of two, **larger groups (3 or more) are not allowed**. If you do work as part of a team, make sure that you indicate this *very clearly* as part of your submission. For homework assignments, each file you submit should have a comment at the top with you and your partner's names and GMU ID numbers. For problem sets, be sure to include both of your names at the top of your submitted PDF. It is important that anything you submit reflects your understanding of the material.

### Late Policy

You have three free late days to be used at your discretion throughout the semester, for either homework or problem set submissions. That means you might turn in one assignment two days late or two different assignments one day late, etc. To avoid an accounting nightmare, you can't pick and choose when to apply late days: a late day is "used" as soon as you submit an assignment late. The first late day is used when you submit one minute after the due date. Two late days are used when you submit 24 hours and one minute after the due date. A third late day is used 48 hours and one minute after the due date, etc. This also means you do not have to let us know ahead of time if you're using a late day. After the free late days are exhausted, you will receive a 20% penalty per day.

### Exams

There will be one exam in this class, a final (cumulative). The date for the final is fixed by the registrar (see the registrar's [final exam schedule](#)). You are allowed one 8.5x11in sheet of notes, front and back. There will be no make up unless previously arranged (*well in advance*).

### Grading policies

Your TA and I will strive to provide you reasonably detailed and timely feedback on every assignment and exam. If you have any questions about any of your grades please reach out to us, either by coming to scheduled office hours or via your "@gmu.edu" email address. If there is an error with your grade, please contact us within a week of when feedback is returned, otherwise we might not be able to change it.

Point breakdown:

- Problem sets: 12% (3% each)
- Homework: 28% (14% each)
- Final: 60%

#### Academic Integrity

All of the assignments you submit in this class must represent your own work. There are absolutely solutions to some of these assignments on the Internet. Do not use them. If you can find them, so can we, and we will check. It's probably a bad idea to even look for them, even for "an idea of how to start" because once you see **one** solution, its hard not to think of it as **the** solution. If you do end up "just looking", make sure you document it in your code. Otherwise, you run the risk of appearing to misrepresent someone else's work as your own. When in doubt, **be explicit** about where the code came from.

### Accommodations

If you need academic accommodations, please make sure you contact the instructor at the beginning of the semester or as soon as possible. Also make sure to contact GMU's Disability Services, available online (<https://ds.gmu.edu>, [ods@gmu.edu](mailto:ods@gmu.edu)) and by phone (703.993.2474), which coordinates all academic accommodations. After you have contacted ODS, you still need to contact the instructor so that appropriate arrangements can be made.

### Topics

The following is the tentative list of topics for the spring 2023 semester, roughly in the order they will be covered. Please check your email and Blackboard regularly for any changes, as this website may not be updated immediately.

#### Week 1

Introductions and Logistics. **Reading:** Ch 1-2

**Week 2 — Coding refresher problem set due (tentative).**

Uninformed search. **Reading:** Ch 3.1-3.4

A\* search and heuristics. **Reading:** Ch 3.5-3.6

**Week 3 — Search problem set due (tentative).**

Search in different domains. **Reading:** Ch 5.1-5.3, 5.5

Probability and state uncertainty. **Reading:** Ch 13

**Week 4 — Search homework due (tentative).**

Bayes Nets, Filtering & Smoothing. **Reading:** Ch 14.1-14.5, 15.1-15.3

Markov Decision Processes. **Reading:** Ch 17.1

**Week 5 — Bayes nets problem set due (tentative)**

Value Iteration and Q-learning. **Reading:** Ch 17.1, Ch 21.3.2

Intro to Machine Learning. **Reading:** Ch 18.1, 18.2, 18.6.1, 18.6.2.

**Week 6 — Reinforcement Learning problem set due (tentative).**

Classification, k-NN, Logistic Regression. **Reading:** Ch 18.6.4, 18.8.1.

Decision Trees/Neural Networks. **Reading:** Ch 18.3/18.7

**Week 7 — Machine Learning homework due (tentative).**

Clustering and k-Means, Gaussian Mixture Models and EM. **Reading:** Ch 20.3

Additional topics, review. **Reading:** none