CS5100: Foundations of Artificial Intelligence (Summer 2026)

Wednesday-Thursday, 3.20 - 5.00 pm

Hurtig Hall, 224, Northeastern University, Boston

1 Status

Welcome to the course! Last updated: 06/28/2026.

2 General Information

In this course, we will work through several of the foundational pillars of AI, including search, sequential decision-making under uncertainty, and both classical and modern machine learning. These subjects, though substantial in their own right, represent only a portion of a much wider body of work that constitutes modern AI, and you should be aware from the outset that certain canonical areas of the discipline fall outside the scope of our semester, notably logical reasoning, constraint satisfaction, and probabilistic graphical models. However, the conceptual and methodological foundation you will develop will prove sufficiently robust to support your serious engagement with a wide class of real-world AI problems, to enable you to learn new techniques on your own as the field advances, and to provide a firm foundation for more specialized graduate coursework and research in AI.

3 Teaching Staff

Role	Name	Email	Office Hours	Location
Instructor	Alesia Chernikova	a.chernikova@northeastern.edu	Wed/Thu 5.00 pm - 6.00 pm	Hurtig Hall 224

You can also schedule an appointment with the instructor. Please email to a.chernikova@northeastern.edu with the subject "CS5100, Appointment" with some time slots that work for you.

4 Content and Communication

We will use Canvas and Piazza as the main tools for course materials and communication:

- Canvas `https://northeastern.instructure.com/courses/256457`: lecture slides, assignments, grades, and other materials. All official course-related content will be posted on Canvas. All assignments should be submitted on Canvas, and all grades will be posted on Canvas.
- Piazza `https://piazza.com/northeastern/summer2026/5100`: announcements, Q&A, and discussion. All communication-related activities will be conducted on Piazza, including official course announcements. Questions and notes may be posted anonymously or with your identity shown, and may also be posted privately to instructors only.

5 Prerequisites

Please familiarize yourself with the following topics:

- Calculus and Linear Algebra
- Probability and Statistics
- Basic Optimization Theory and Differential Calculus
- Algorithms and Data Structures
- Python programming language

6 Work Load and Policies

6.1 Homework Assignments

There will be 5 homework assignments. They will contain both theoretical questions and programming blocks. Students may discuss the problems with other students but must write up their own solutions. In each assignment, please also indicate with whom you discussed the problems (if anyone). You may engage in discussion with AI tools, but you should not ask them to solve assignments, and you should not copy answers/code from them. Additionally, if you use these AI tools, you should acknowledge their usage and provide details about how you are using them. Lateness: Up to 1 day late (24 hours), penalized by 10%. Submissions more than 24 hours late will not be accepted.

6.2 Research Project

There will be 1 research project. Students may select any topic relevant to AI. Some examples will be provided later in the class. The project may be done in a group of 2 or 3 people. One-person groups are not encouraged, though case-by-case exceptions may be made. You may engage in discussion with AI tools, but you should acknowledge their usage and provide details about how you are using them. Lateness: Up to 1 day late (24 hours), penalized by 10%. Submissions more than 24 hours late will not be accepted. The project has the following milestones:

- Project pre-proposal
- Project proposal
- Preliminary results
- Project report draft
- Project report

6.3 Exams

There will be 2 midterm exams. There will be no final exam. No collaboration, textbook, or AI usage is allowed during exams. You may bring two single-sided A4 pages of handwritten or printed notes, which must be turned in with the completed exam.

7 Academic Integrity

Cheating and other acts of academic dishonesty will be referred to OSCCR (office of student conduct and conflict resolution) `https://osccr.sites.northeastern.edu` and the Khoury College of Computer Sciences.

8 Grades

The grades will be based on the following:

- Homework assignments - 40%
- Research project - 30%
- Exam1 + Exam2 - 25%
- Class participation - 5%

9 Textbooks and References

There is no required textbook, but you might find the following resources helpful:

- Artificial Intelligence: A Modern Approach (AIMA), by Stuart Russell and Peter Norvig. `https://aima.cs.berkeley.edu`. The link takes you to the 4th edition, which is worth considering buying if you plan to continue working in AI. The 3rd edition can be found online and is reasonably similar to the current 4th edition.
- Understanding Machine Learning: From Theory to Algorithms, by Shai Ben-David, and Shai Shalev-Shwartz. `https://www.cs.huji.ac.il/~shais/UnderstandingMachineLearning/understanding-machine-learning-theory-algorithms.pdf`. This book is highly recommended for those who want to learn more about the theoretical side of machine learning.

- Pattern Recognition and Machine Learning, by Christopher Bishop. https://www.microsoft.com/en-us/research/wp-content/uploads/2006/01/Bishop-Pattern-Recognition-and-Machine-Learning-2006.pdf. This book provides a Bayesian-focused introduction to AI.
- An Introduction to Statistical Learning: with Applications in Python, by Gareth James, Daniela Witten, Trevor Hastie, Robert Tibshirani, and Jonathan Taylor. https://www.statlearning.com. This book covers the fundamentals of machine learning.
- Mathematics for Machine Learning, by A. Aldo Faisal, Cheng Soon Ong, and Marc Peter Deisenroth. https://mml-book.github.io/book/mml-book.pdf. This book covers some mathematical fundamentals needed to study AI, such as Linear Algebra, Analytic Geometry, Matrix Analysis, Probability, and Optimization.
- Berkley, CS 188, Spring 26, developed by Dan Klein, Stuart Russell, and many others. https://inst.eecs.berkeley.edu/~cs188/sp26/. The link has many useful resources, such as video recordings and practice worksheets with solutions.
- Probability Review, by David Blei. https://www.cs.princeton.edu/courses/archive/spring07/cos424/scribe_notes/0208.pdf.
- Probability Theory Review, by Arian Maleki and Tom Do. https://cs229.stanford.edu/section/cs229-prob.pdf.
- Linear Algebra Review and Reference, by Zico Kolter. https://cs229.stanford.edu/section/cs229-linalg.pdf
- Optimization for Machine Learning, by Elad Hazan. https://arxiv.org/pdf/1909.03550.
- Python 3 documentation. https://docs.python.org/3/tutorial/index.html.
- Github documentation. https://docs.github.com/en/get-started/start-your-journey.
- AI courses, by Sebastian Rashka. https://sebastianraschka.com/teaching/. This link provides some useful resources on modern deep learning.
- Some interesting thoughts from Nicolas Carlini. https://nicholas.carlini.com/writing.
- Anthropic AI safety research. https://www.anthropic.com/research.
- AI security institute. https://www.aisi.gov.uk.

10 Schedule

Subject to change.

Date	Lecture	Topic	Resources	Assignments due on 11.59 pm
05/06	1	Introduction, Agents, Environments	AIMA 1–2	
05/07	2	Uninformed search: BFS, UCS	AIMA 3.1–3.4	
05/13	3	Uninformed search	AIMA 3.1–3.4	
05/14	4	Heuristic search	AIMA 3.5–3.6	
05/20	5	Heuristic search	AIMA 3.5–3.6	
05/21	6	Adversarial search	AIMA 5.1-5.5	
05/27	7	Adversarial search	AIMA 5.1-5.5	
05/28	8	Decision Theory	AIMA 16.1-16.3	
05/31				HW1
06/03	9	Markov Decision Processes	AIMA 17.1-17.2	
06/04	10	Markov Decision Processes	AIMA 17.1-17.2	
06/07				Project pre-proposal
06/10	11	Reinforcement Learning	AIMA 17.3, 22.1-22.2	
06/11	12	Reinforcement Learning	AIMA 22.2-22.4, 22.7.1	
06/14				
06/17		Exam 1	Lectures 2-7	
06/18	13	Intro to Machine Learning, Linear Regression	AIMA 19.19-19.2, 19.6 -19.6.2	
06/21				HW2
06/24	14	Logistic regressions	AIMA 19.6.4-19.6.5	
06/25	15	Deep Learning, backpropagation	AIMA 21.1-21.2, 21.4	
06/28				Project Proposal
07/01	16	CNNs and sequence models	CNNs, CNNs,Sequence models	
07/02	17	Transformers	T1 T2 T3 LSTMs	
07/05				
07/08		Invited lecture, Germans Savcisens	life2vec	
07/09	18	Unsupervised ML		
07/12				HW3
07/15	19	Deep RL		
07/16	20	Generative AI		
07/19				Project Milestone
07/22	21	LLMs and Frontiers		
07/23		Exam 2	Lectures 9-19	
07/26				HW4
07/29	22	Responsible AI		
07/30		Invited lecture, Konstantina Bairaktari		
08/02				Project report draft
08/05		Topics of interest/ Research		
08/06		Projects presentations		
08/09				Project report