COMSW4701_001_2020_1 - ARTIFICIAL INTELLIGENCE

Jump to Today ✎ Edit

Welcome to Artificial Intelligence!

THIS PAGE IS ALL YOU NEED TO GET STARTED

In this course, you will learn the fundamental concepts of Artificial Intelligence (AI) and apply them to the design and implementation of intelligent agents that solve real-world AI problems, including problems in search, games, machine learning, logic, and constraint satisfaction.

We will provide a broad understanding of the basic techniques for building intelligent computer systems. Topics include the history of AI, intelligent agents, state-space problem representations, uninformed and heuristic search, game playing and adversarial search, logical agents, constraint satisfaction problems, along with techniques in machine learning and other applications of AI, such as natural language processing (NLP).

Time: Tuesday and Thursdays @10:10am-11:25am

Room: 501 Northwest Corner Building

Piazza: piazza.com/columbia/spring2020/comsw4701_001_2020_1artificialintelligence (https://piazza.com/class/k5khw6319ij4vf) (https://piazza.com/class/k5jz39wa41067d)

Gradescope code: TBA

Instructor:

- Ansaf Salleb-Aouissi ansaf@cs.columbia.edu (mailto:ansaf@cs.columbia.edu)
- Office hours: TBA. Location: 702 CEPSR.

Office hours will start the week of MONDAY JANUARY 27.

Teaching Assistants:

TA pictures will be available in the TA_PICS folder.

1. Nicole Mbithe (Head TA):
2. Nicolae Lari
3. Uzay Macar
4. Nuneke Kwetey
5. Nic Cornejo
6. Robert Costales
7. Anastasia Dmitrienko
8. Tanvi Hisaria
9. Daniel Jaroslawicz
10. Beom Joon Baek

Office hours schedule:

TA Office Hours below. Some TAs may need to change their OH during the semester. TA office hours will be held in the TA room in Mudd 122A or Barnard CS TA room (502 Milstein).

	Monday	Tuesday	Wednesday	Thursday	Friday
Morning					

Afternoon					

Pre-requisites:

You will need some basics in Linear algebra (vectors, matrices, derivatives), Calculus, Probability theory.

You will need to be fluent in programming in Python 3 programming and know data structures.

Course Level:

Expect to spend several hours to complete the assignments, although the exact amount of time will depend on your background and proficiency with coding. So start your homework early.

Assignments:

There will be about 6 homework assignments, each consisting of a conceptual and programming portion. Students will have about two weeks to complete each assignment. Assignments have to be done individually, no collaboration/group work. **Please submit your *own* work.**

Grading: Assignments: 50%, Midterm 20%, Final: 30%.

Academic Honesty Policy:

Please familiarize yourself with the honesty, integrity, and **cheating policies of the University and the Department of Computer Science (http://www.cs.columbia.edu/education/honesty)** . Ignorance of these policies will not be considered as an acceptable reason for violating the policies. It is your responsibility to be aware of the policies. Each student is sole owner of his own code and work and **must NOT**:

- Submit work that is not original.
- Publish code or solutions online.
- Post the course questions on forums including stack overflow.
- Submit someone else's work, or a modification of that work, with or without that person's knowledge.
- Allow someone else to submit his/her work, or a modification of that work.
- Solve as a group a quiz or project. **All coursework is to be done by the student working alone.**
- Contract course work out to others.
- Plan or execute with another student a cooperative subterfuge during an exam.
- Make use of unauthorized material during an exam.

Project assignments will be checked with plagiarism detection software.

Thank you for abiding by these rules. Doing so will ensure the experience is fair to everyone taking this class or the future sessions of this class.

Recommended Textbook:

Artificial Intelligence, A Modern Approach. Stuart Russell and Peter Norvig. Third Edition. Pearson Education. Check out the book resources: **http://aima.cs.berkeley.edu/ (http://aima.cs.berkeley.edu/)**

Late Policy:

Homework must be submitted by the deadline**.** Late policy will be available soon.

No make-up exams:

In case you must miss an exam for a well-justified valid medical/family emergency (a simple vague note from a doctor will not be accepted), your grade composition will be adjusted.

Exceptions:

To be considered for an exception to the above policies you must furnish ***both*** a letter from your physician (doctor) describing a medically-necessary delay in your studies **and** a letter sent directly to your instructor from your academic dean. There is no make-up exam.

Regrade requests:

All regrade requests for exams and HW are to be submitted within one week after grades are released.

Disability:

All ODS/disability forms are signed once with the disability office. You no longer need to bring in forms to sign.

Topics

1. Introduction to AI, history of AI, course logistics, and roadmap
2. Intelligent agents, uninformed search
3. Heuristic search, greedy search, A* algorithm, stochastic search
4. Adversarial search, game playing
5. Machine Learning: basic concepts, linear models, K nearest neighbors, overfitting
6. Machine Learning: perceptrons, neural networks, naive Bayes, decision trees, ensemble, logistic regression, and unsupervised learning
7. Constraint satisfaction problems
8. Reinforcement learning.
9. Logical agents, propositional logic and first order logic
10. AI applications to natural language processing (NLP)
11. Review and conclusion

Tentative Schedule:

SUBJECT TO CHANGE. PLEASE CONSULT REGULARLY FOR HW DEADLINES.

Week 1 (21-Jan & 23-jan)

Introduction: history of AI. A quick tour of AI, course overview, logistics

Intelligent agents

Search Algorithms: search agents

Reading/notes:

Chapter 1, 2

Computing Machinery and Intelligence by Alan Turing (1950)

President Obama's Preparing for the Future of Artificial Intelligence.

Artificial Intelligence and Life in 2030

Week 2 (Jan 28 & Jan 30)

Uninformed search

Informed search

Reading/notes:

Chapter 3, 4

Depth-First Iterative Deepening Korf 1985

A formal basis for the heuristic determination of minimum cost paths Hart et al. 1968.

Week 3 (Feb 4 & Feb 6)

Stochastic/Local search

Constraint Satisfaction Problems

Reading/notes:

Chapter 3, 4

An overview of Genetic Algorithms Part I, An overview of Genetic Algorithms Part II

Updated Chapter 6

Week 4 (Feb 11 & Feb 13)

Constraint Satisfaction Problems

Adversarial search. Stochastic games

Reading/notes:

Chapter 5

Programing a computer for playing Chess. Claude Shannon 1950

Checkers is solved. Schaeffer, 2007

Demo: Alpha-Beta pruning

Week 5 (Feb 18 & Feb 20)

Adversarial search. Stochastic games

Machine Learning: Basic Concepts, KNN. Curse of dimensionality.

Reading/notes:

Chapter 18

A Course in Machine Learning by Hal Daumé III

Deep Learning by Goodfellow, Bengio and Courville. The book includes a review of Linear Algebra and Probability and Information Theory

Are ML and Statistics Complementary?

Data Mining and statistics: what's the connection?

When is the NN meaningful?

Week 6 (Feb 25 & Feb 27)

Decision trees

Machine Learning: Naive Bayes

Linear methods (linear regression)

Reading/notes:

Chapter 18

A Course in Machine Learning by Hal Daumé III

Deep Learning by Goodfellow, Bengio and Courville. The book includes a review of Linear Algebra and Probability and Information Theory

Are ML and Statistics Complementary?

Data Mining and statistics: what's the connection?

Week 7 (Mar 3 & Mar 5)

Perceptrons, Neural networks

Perceptron: Proof of convergence

Reading/notes:

Deep Learning by Goodfellow, Bengio and Courville

Tensorflow playground

An Empirical Comparison of Supervised Learning Algorithms Rich Caruana and Alexandru Niculescu-Mizil. ICML 2006

Week 8 (Mar 10 & Mar 12)

Midterm (Mar 12)

Covers Weeks 1-6.

Closed book but one cheat sheet two sides allowed. Calculators allowed but no other electronics.

Ensemble methods

Clustering

Reading/notes:

Some visualizations of clustering

https://www.naftaliharris.com/blog/visualizing-dbscan-clustering/

https://www.naftaliharris.com/blog/visualizing-k-means-clustering/

Week 9 (Mar 16- 20)

SPRING BREAK!

Week 10 (Mar 24 & Mar 26)

Logical Agents

Introduction to NLP

Reading/notes:

Chapter 7

McCarthy: programs with common sense

Week 11 (Mar 31 & Apr 2)

Invited speaker I: Natural Language Processing. Yassine Benajiba

Bayes nets and HMMs

Reading/notes:

Chapter 22

"AI's Language Problem", by Will Knight,
MIT Technology Review, August 9, 2016.

https://www.technologyreview.com/s/602094/ais-language-problem/

Week 12 (Apr 7 & Apr 9)

Invited speaker II: Introduction to Computer Vision. Carl Vondrick (April 7)

Bayes nets and HMMs

Reading/notes:

Reading: A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition

---

Week 13 (Apr 14 & Apr 16)

MDPs and RL (TBA)

Reading/notes:

Week 14 (Apr 21 & Apr 23)

Fairness in AI (TBA)

Interpretable AI

Reading/notes:

Week 15 (Apr 28 and April 30)

AI and ethics, AI and Inclusion, Conclusion

Reading/notes:

Artificial intelligence: How to avoid racist algorithms

Joy Buolamwini How I'm fighting bias in algorithms

David Lee: Why jobs of the future won't feel like work

DURING FINAL WEEK:

Date according to **registrar (https://ssol.columbia.edu/cgi-bin/ssol/CC5FSo3KqsZ4HAfHAmBs1h/?p%.5Fr%.5Fid=CC5FSo3KqsZ4HAfHAmBs1h&p%.5Ft%.5Fid=1&tran%.5B1%.5D%.5Fterm%.5Fid=20201&tran%.5B1%.5D%.5Ftran%.5Fname=spex)** schedule:

Thursday May 14 @9am-12pm. (duration will be confirmed).

FINAL (Non cumulative)

Covers Weeks 7-14.

Closed book but one cheat sheet two sides allowed. Calculators allowed but no other electronics.

Course Summary:

Date	Details