Agoura Engineering Circle Introduction to AI Curriculum

Introduction to Artificial Intelligence Using Python

Computers have been an integral part of our lives in the past few decades. Most of the applications, we use, have been "programmed" to do what they do. For example, take your iPhone or Android phone, the operating system behind your phone, is (for the most part) explicitly programmed to do what it does. However, in the last decade or two, we've been seeing a rise, in the kind of programs which don't have to be programmed to do their job. They instead "learn" or behave "intelligently". You may have heard of one or more of these: Chat bots, high frequency trading (buying and selling of stocks), crop weed detection, detecting patients with heart disease, Alpha Go (the program that defeated the best Go player!) or Alpha Zero (the chess program that learnt chess in one day) and self-driving cars. While these applications have been programmed, the programming is not the traditional kind where they program exactly what to do. They program a learning mechanism and it learns how to play or respond to the unique situations that arise during the course of running the program.

Imagine building a self-driving car and having to program exactly how to turn right at every right turn in the world. Sometimes, there may be people around the corner, it may be raining or another car might be close behind. Programming all these situations explicitly is almost impossible. Instead these programs learn based on previously known "good" states and adapt to the new situations.

In this course, we'll learn to program an intelligent application, specifically, predicting the success of a movie. We'll do this as a four-part, 90-hour course consisting of 36 sessions:

1. **Mathematical Foundations for AI using NumPy**
2. **Introduction to Artificial Intelligence methods**

Project: Predicting success of a movie.

Course flow

We'll start out our journey by introducing basic programming concepts like branching, iteration, modular coding and data structures while solving mathematics problems that most students can understand but not necessarily solve by hand easily. This type of interdisciplinary learning helps in learning two things at the same time along with reinforcing any prior knowledge.

Once everyone has achieved some proficiency in programming, we'll move on to solving artificial intelligence and machine learning challenges.

Class Schedule

Alternate Saturdays 9:30am – 12:00pm (online).

Class	Homework Review	Quiz	Lecture & Class Work
Intro to AI	9: 30 – 10 AM	10- 10: 15 AM	10: 15 AM -12: 00 PM

Please note that if you miss two classes in a row, you'll not be able to follow the topics anymore. Students must login 10 min before start of the class.

Awards

At the end of each Semester (Fall, Spring), 3 students will receive the Star Award which is based on the cumulative score from quiz, home work and Final Exam.

Course Details

Pre-requisites: Good understanding of basic mathematical concepts (no higher than 8th grade level or taught in Senior Intermediate level at Agoura Math Circle). Also, students have to get their own desktops or laptops to program and **completed Introduction Data Science with good grade**.

Course Registration & Website: Each student should register for the class using their own email ID (Not parent's email ID). All course communication, homework submission will be through the course website. Each student has to register for each one of the four sessions irrespective of whether they have previously registered or not. Preference will be given to students who have attended previous sessions.

Class workload: Apart from the 2.5 hours of class once every 2 weeks, students are expected to spend at least 1 hour every day of the week for a total work load of between 15 – 20 hours. If you can't make this commitment, please do not register. The course material to be covered is pretty heavy and if you fall behind, catching up is difficult.

Class style: Classes will be Skype based in an interactive manner involving discussions and coding either individually or as a team. Instead of striving towards finishing certain amount of material in each class, we'll work towards certain milestones which involve writing a few programs individually or as a team. During the courses, a textbook will be recommended for each part.

Homework Policy: After each class, homework assignments will be mailed out. They are due 11 days from that day i.e. Wednesday of the next week. Most of the homework problems are fairly challenging, especially to those without any previous experience in programming. Please feel free to discuss homework problems with friends, parents or anyone else. But the final submission should be yours. Any sort of plagiarism will not be tolerated.

Final Exam Policy: Each semester has a final exam worth 100 points. The test will be online for 3 hours.

What will you learn from this course?

1. Achieve decent proficiency in programming with Python
2. Setup a GitHub portfolio to show case your work
3. Improvement in report writing and presentation skills.
4. Team work and collaboration towards finishing a project in artificial intelligence using open-source libraries.
5. Increase in confidence to tackle problems in a logical and algorithmic fashion.

In order to come up with the list of topics in a manner that is comprehensive and meets our objectives, a few online resources including but not limited to Coursera, pandas have been used. While they were used to identify some topics, the material will be created by the developers of the course.

Syllabus

Fall Semester

Mathematical Foundations for AI using NumPy

Module 1: Linear Algebra

1. Session 1: Linear Equations & Matrices
    a. Introduction to linear equations
    b. Solving linear equations using matrices
    c. Introduction to NumPy
    d. Matrix multiplication
2. Session 2: Higher dimensional matrices:
    a. Introduction to higher dimensional matrices aka Tensors
    b. Introduction to Tensor flow

Module 2: Probability

3. Session 3: Introduction to Probability
    a. Definition of probability, Bayes Theorem and Independent Events
    b. Programming probabilistic problems
    c. Introduction to Combinatorics

4. Session 4: Probability continued
    a. Combinatorics continued
    b. Introduction to data distributions: Cumulative distribution function
    c. Exploring probabilistic distributions
    d. Plotting Histograms

Module 3: Statistics

5. Session 5: Introduction to Statistics
    a. Introduction to statistics
    b. Mean, Median, Standard Deviation
    c. Programming statistical problems
6. Session 6: Hypothesis testing
    a. Statistical testing
    b. Introduction to hypothesis testing
    c. Outliers
    d. Pitfalls of hypothesis testing

Module 4: Data Visualization

7. Session 7: Charting
    a. Introduction to charting
    b. Basic charting explored

8. Session 8: Advanced charting
    a. Exploring advanced charting options

9. Session 9: Projects

10. Final Exam

Spring Semester

Introduction to Artificial Intelligence methods

Session 1: Introduction to machine learning and artificial intelligence

a. What is machine learning?
b. What is artificial intelligence
c. Various keywords explained
d. Introduction to supervised and unsupervised learning

Module 1: Supervised Machine Learning

Session 2: Classification problem

e. Introduction to classification
f. K-Nearest neighbor

Session 3: Linear Regression

g. Introduction to linear regression
h. Discussion about cost function
i. Overfitting and under-fitting
j. Bias and variance

Session 4: Regularization

k. Introduction to regularization
l. Lasso and Linear Ridge regularization
m. Polynomial regularization

Session 5: Machine learning practices:

n. Feature selection
o. Testing, training and validation

Session 6: Logistic regression

: Classification using Support Vector Machines

Module 2: Model evaluation and Selection

Session 7: Model evaluation

p. How do we evaluate the performance of a model?
q. Confusion Matrices
r. Evaluation metrics

Session 8: Model evaluation continued

s. Regression Evaluation
t. Optimizing Classifiers for Different Evaluation Metrics

Session 9: Advanced Topics

a. Naive Bayes Classifiers
b. Session 11: Random Forests

Session 10: Project: Predicting success of a movie

1. Getting data:
    i. Get data from IMDB
    ii. Get data from the-Numbers.com
2. Identifying features
3. Predicting movie success as a Regression problem
    i. Using linear regression
    ii. Using random forests

Predicting movie success as a classification