ENTERPRISE

UDACITY

SCHOOL OF ARTIFICIAL INTELLIGENCE

AI Programming with Python

Nanodegree Program Syllabus

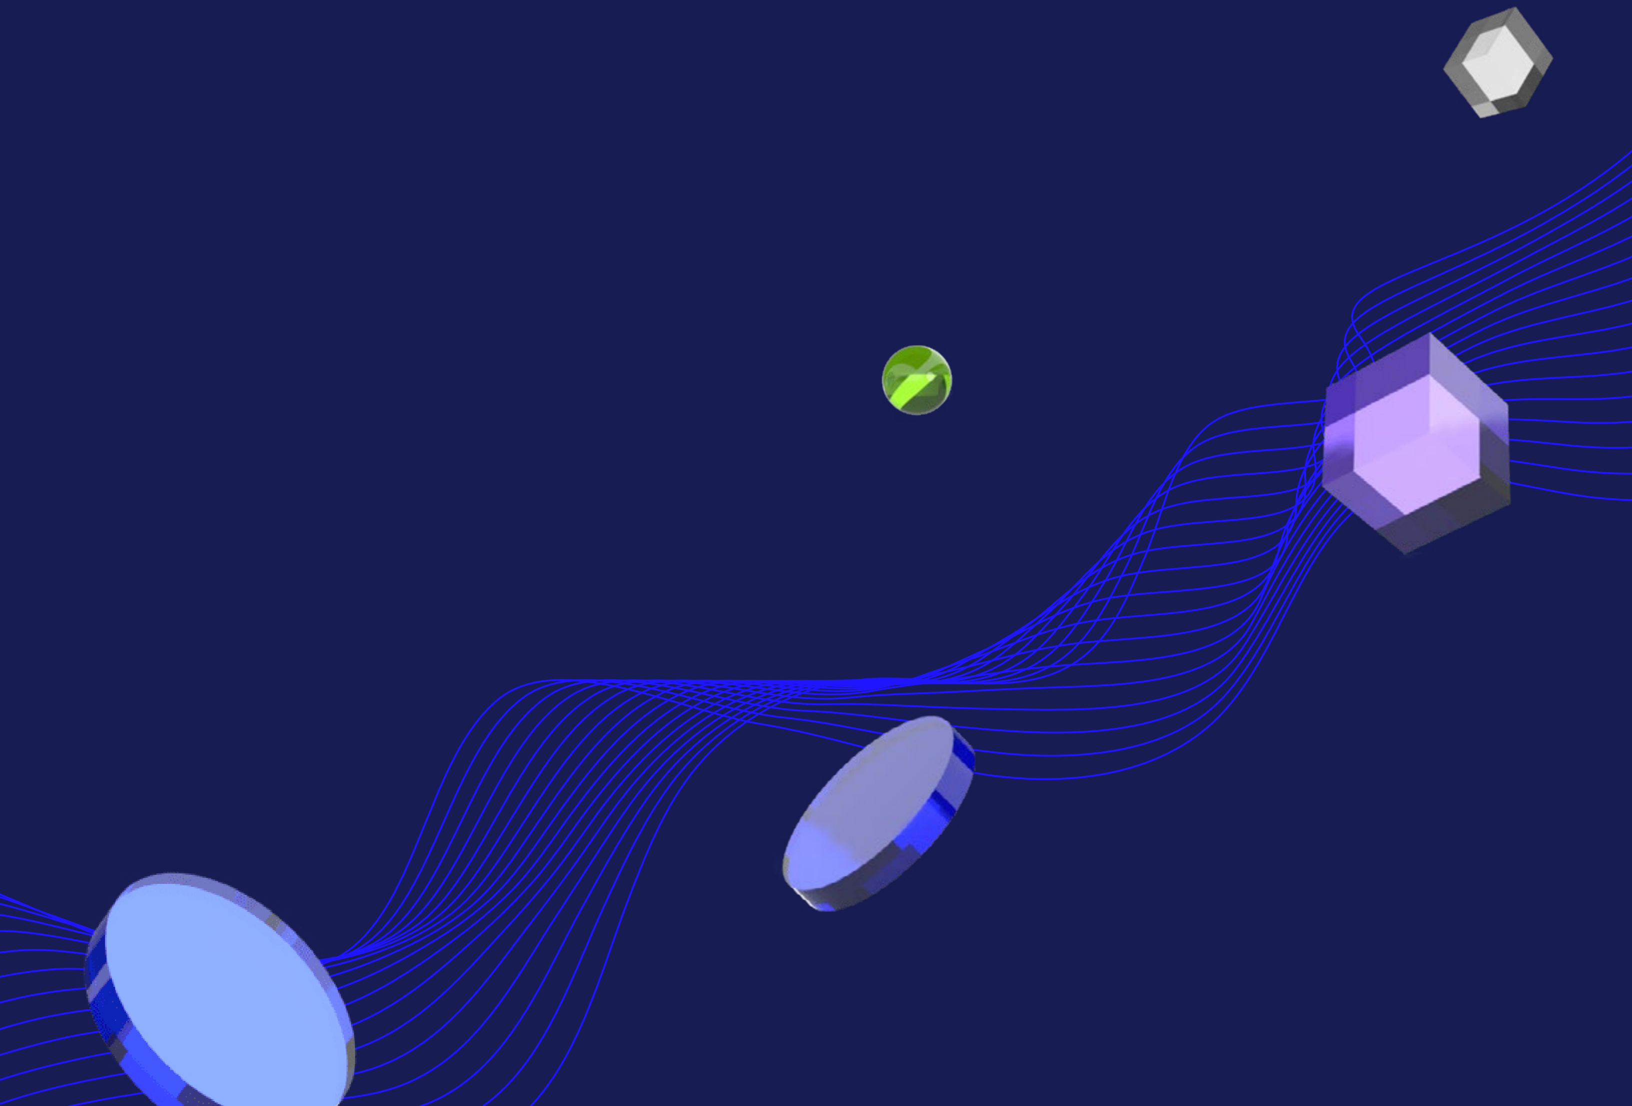

Overview

Start using AI techniques and developing skills related to programming, linear algebra, and neural networks.

Learning Objectives

A graduate of this program will be able to:

- Utilize programming tools like Python, NumPy, and PyTorch
- Understand Calculus and linear algebra applications
- Perform key techniques of neural networks, including gradient descent and backpropagation

Program information

Estimated Time

3 months at 10hrs/week*

Skill Level

Beginner

Prerequisites

Learner should have basic calculus knowledge, including how to calculate derivatives. Basic algebra and programming knowledge will shorten the time to mastery.

Required Hardware/Software

Learners need access to a computer running OS X or Windows.

*The length of this program is an estimation of total hours the average student may take to complete all required coursework, including lecture and project time. If you spend about 5-10 hours per week working through the program, you should finish within the time provided. Actual hours may vary.

Course 1

Introduction to Python

Start coding with Python, drawing upon libraries and automation scripts to solve complex problems quickly.

Course Project

Use a Pre-trained Image Classifier to Identify Dog Breeds

In this project, learners will be testing their newly-acquired Python coding skills by using a trained image classifier. They will need to use the trained neural network to classify images of dogs (by breeds) and compare the output with the known dog breed classification. Learners will have a chance to build their own functions, use command line arguments, test the runtime of the code, create a dictionary of lists, and more.

Lesson 1

Why Python?

- Learn why we program.
- Prepare for the course ahead with a detailed topic overview.
- Understand how programming in Python is unique.

Lesson 2

Data Types & Operators

- Understand how data types and operators are the building blocks for programming in Python.
- Use the following data types: integers, floats, booleans, strings, lists, tuples, sets, and dictionaries.
- Use the following operators: arithmetic, assignment, comparison, logical, membership, and identity.

Lesson 3

Control Flow

- Implement decision-making in your code with conditionals.
- Repeat code with for and while loops.
- Exit a loop with break, and skip an iteration of a loop with continue.
- Use helpful built-in functions like zip and enumerate.
- Construct lists in a natural way with list comprehensions..

Lesson 4

Functions

- Write your own functions to encapsulate a series of commands.
- Understand variable scope, i.e., which parts of a program variables can be referenced from.
- Make functions easier to use with proper documentation.
- Use lambda expressions, iterators, and generators.

Lesson 5

Scripting

- Write and run scripts locally on your computer.
- Work with raw input from users.
- Read and write files, handle errors, and import local scripts.
- Use modules from the Python standard library and from third-party libraries.
- Use online resources to help solve problems.

Lesson 6

Classes

- Object-oriented programming provides a few benefits over procedural programming. Learn the basics by understanding how to use Classes.

Course 2

Anaconda, Jupyter Notebook, NumPy, Pandas & Matplotlib

Learn how to use all the key tools for working with data in Python: Jupyter Notebooks, NumPy, Anaconda, Pandas, and Matplotlib.

Lesson 1

Anaconda

- Learn how to use Anaconda to manage packages and environments for use with Python.

Lesson 2

Jupyter Notebooks

- Learn how to use Jupyter Notebooks to create documents combining code, text, images, and more.

Lesson 3

Numpy Basics

- Learn the value of NumPy and how to use it to manipulate data for AI problems.
- Mini-Project: Use NumPy to mean normalize an darray and separate it into several smaller ndarrays.

Lesson 4

Panda Basics

- Learn to use Pandas to load and process data for machine learning problems.
- Mini-Project: Use Pandas to plot and get statistics from stock data.

Lesson 5

Matplotlib Basics

- Learn how to use Matplotlib to choose appropriate plots for one and two variables based on the types of data you have.

Course 3

Linear Algebra Essentials

Learn the foundational math needed for AI success—vectors, linear transformations, and matrices—as well as the linear algebra behind neural networks.

Lesson 1

Introduction

- Learn the basics of the beautiful world of linear algebra and why it is such an important mathematical tool.

Lesson 2

Vectors

- Learn about the basic building block of linear algebra.

Lesson 3

Linear Combination

- Learn how to scale and add vectors and how to visualize them in 2 and 3 dimensions.

Lesson 4

Linear Transformation & Metrices

- Learn what a linear transformation is and how is it directly related to matrices. Learn how to apply the math and visualize the concept.

Lesson 5

Linear Algebra in Neural Networks

- Learn about the world of neural networks and see how it related directly to neural networks.

Lesson 6

Labs

- Vectors Lab: Learn howto graph 2D and 3D vectors.
- Linear Combinations Lab: Learn how to computationally determine a vector’s span and solve a simple system of equations.
- Linear Mapping Lab: Learn how to solve problems computationally using vectors and matrices.

Course 4

Calculus Essentials

Learn the foundations of calculus to understand how to train a neural network: plotting, derivatives, the chain rule, and more. See how these mathematical skills visually come to life with a neural network example.

Lesson 1

Introduction

- Visualize the essence of calculus. Learn why it is such a powerful concept in mathematics.

Lesson 2

Derivates Through Geometry

- Learn about the derivative, one of the most important tools in calculus.
- See how a derivative can measure the steepness of a function and why it is such an important indicator in the world of machine learning.

Lesson 3

Chain Rule & Dot Product

- Learn how to find the derivative of a composition of two or more functions, a very important tool in training a neural network.

Lesson 4

More on Derivatives

- Learn more about derivatives while focusing on exponential and implicit functions.

Lesson 5

Limits

- Learn about the formal definition of a derivative through understanding limits.

Lesson 6

Integration

- Learn about the inverse of a derivative: the integral.

Lesson 6

Calculus in Neural Networks

- Learn more about the world of neural networks and see how it relates directly to calculus through an explicit example.

Course 5

Neural Networks

Gain a solid foundation in the latest trends in AI: neural networks, deep learning, and PyTorch.

Course Project

Create Your Own Image Classifier

Successful software developers need to know how to incorporate deep learning models into everyday applications. Any device with a camera will be using image classification, object detection, and face recognition, all based on deep learning models. In this project learners will implement an image classification application. This application will train a deep learning model on a dataset of images. It will then use the trained model to classify new images. First learners will develop their code in a Jupyter notebook to ensure their training implementation works well. Then, they will convert their code into a Python application that they will run from the command line of their system.

Lesson 1

Introduction to Neural Networks

- Acquire a solid foundation in deep learning and neural networks.
- Implement gradient descent and backpropagation in Python.

Lesson 2

Training Neural Networks

- Learn about techniques for how to improve training of a neural network, such as early stopping, regularization and dropout..

Lesson 3

Deep Learning with PyTorch

- Learn how to use PyTorch for building deep learning models.

Meet your instructors.

Ortal Arel

Instructor

Ortal Arel holds a PhD in computer engineering, and has been a professor and researcher in the field of applied cryptography. She has worked on design and analysis of intelligent algorithms for high-speed custom digital architectures.

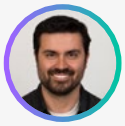

Luis Serrano

Machine Learning Engineer

Luis was formerly a machine learning engineer at Google. He holds a PhD in mathematics from the University of Michigan, and a postdoctoral fellowship at the University of Quebec at Montreal.

Juan Delgado

Computational Physicist

Juan is a computational physicist with a master's degree in astronomy. He is finishing his PhD in biophysics. He previously worked at NASA developing space instruments and writing software to analyze large amounts of scientific data using machine learning techniques.

Mat Leonard

Instructor

Mat is a former physicist, research neuroscientist, and data scientist. He completed his PhD and postdoctoral fellowship at the University of California, Berkeley.

Mike Yi

Content Developer

Mike is a content developerwith a BS in mathematics and statistics. He received his PhD in cognitive science from the University of Irvine. Previously, he worked on Udacity's Data Analyst Nanodegree program as a support lead.

Grant Sanderson

Instructor

Grant Sanderson is the creator of the YouTube channel 3Blue1Brown, which is devoted to teaching math visually, using a custom-built animation tool. He was previously a content creator for Khan Academy.

Jennifer Staab

Statistician & Computer Scientist

Jennifer holds a PhD in computer science and a master's in biostatistics; she was a professor at Florida Polytechnic University. She previously worked at RTI International and United Therapeutics as a statistician and computer scientist.

Juno Lee

Data Scientist at Looplist

As a data scientist at Looplist, Juno built neural networks to analyze and categorize product images, a recommendation system to personalize shopping experiences for each user, and tools to generate insight into user behavior.

Andrew Paster

Entrepreneur

Andrew has an engineering degree from Yale, and has used his data science skills to build a jewelry business from the ground up. He has additionally created courses for Udacity's Self-Driving Car Engineer Nanodegree program.

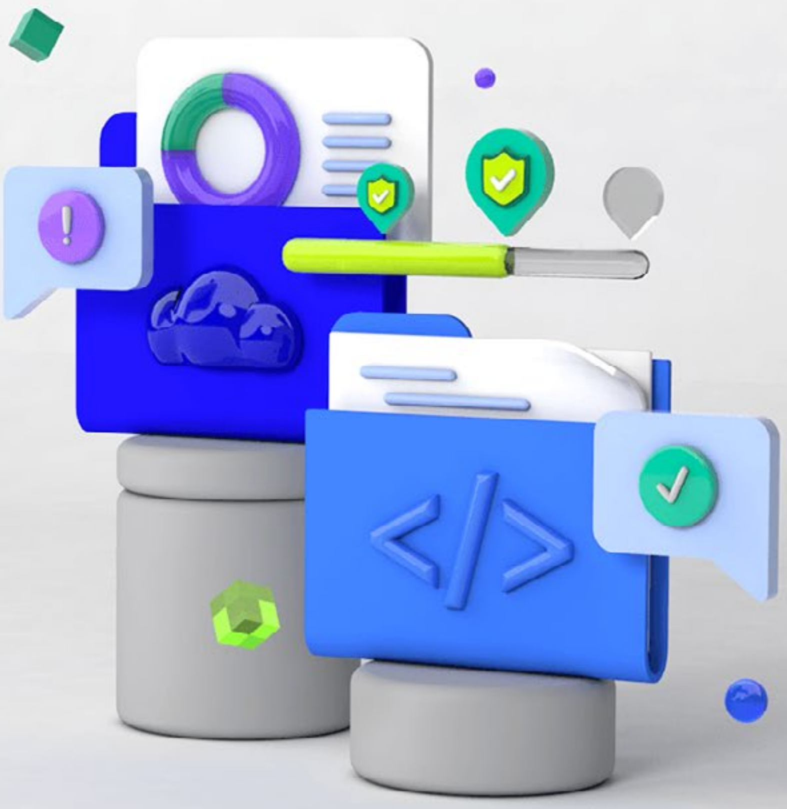

Udacity's learning experience

Hands-on Projects

Open-ended, experiential projects are designed to reflect actual workplace challenges. They aren't just multiple choice questions or step-by-step guides, but instead require critical thinking.

Quizzes

Auto-graded quizzes strengthen comprehension. Learners can return to lessons at any time during the course to refresh concepts.

Knowledge

Find answers to your questions with Knowledge, our proprietary wiki. Search questions asked by other students, connect with technical mentors, and discover how to solve the challenges that you encounter.

Custom Study Plans

Create a personalized study plan that fits your individual needs. Utilize this plan to keep track of movement toward your overall goal.

Workspaces

See your code in action. Check the output and quality of your code by running it on interactive workspaces that are integrated into the platform.

Progress Tracker

Take advantage of milestone reminders to stay on schedule and complete your program.

Our proven approach for building job-ready digital skills.

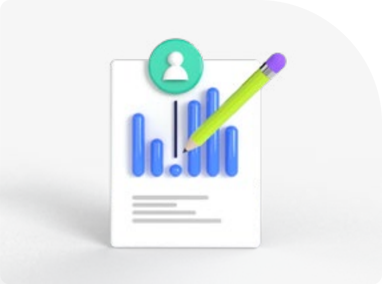

Pre-Assessments

Identify skills gaps.

- In-depth assessments benchmark your team's current level of knowledge in key areas.
- Results are used to generate custom learning paths.

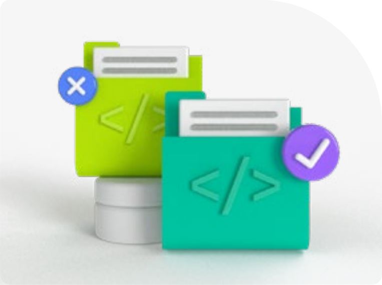

Experienced Project Reviewers

Verify skills mastery.

- Personalized project feedback and critique includes line-by-line code review from skilled practitioners with an average turnaround time of 1.1 hours.
- Project review cycle creates a feedback loop with multiple opportunities for improvement—until the concept is mastered.
- Project reviewers leverage industry best practices and provide pro tips.

Technical Mentor Support

24/7 support unblocks learning.

- Learning accelerates as skilled mentors identify areas of achievement and potential for growth.
- Unlimited access to mentors means help arrives when it's needed most.
- 2 hr or less average question response time assures that skills development stays on track.

Mentor Network

Highly vetted for effectiveness.

- Mentors must complete a 5-step hiring process to join Udacity's selective network.
- After passing an objective and situational assessment, mentors must demonstrate communication and behavioral fit for a mentorship role.
- Mentors work across more than 30 different industries and often complete a Nanodegree program themselves.

Dashboard & Reporting

Track course progress.

- Udacity's enterprise management console simplifies management of bulk enrollments and employee onboarding.
- Interactive views help achieve targeted results to increase retention and productivity.
- Maximize ROI while optimizing job readiness.

Learn more at

udacity.com/enterprise →

12.02.22 | V1.0