



NANODEGREE PROGRAM SYLLABUS

Machine Learning Engineer with Microsoft Azure



Overview

In this program, students will enhance their skills by building and deploying sophisticated Machine Learning (ML) solutions using popular open source tools and frameworks such as scikit-learn. Using Azure Machine Learning's MLOps capabilities, students will gain experience in understanding their ML models, protecting people and their data, and controlling the end-to-end ML lifecycle at scale. Gain practical experience by using the built-in Azure labs accessible inside the Udacity classroom and run complex machine learning tasks for no additional cost.

Prerequisites:

- Experience with basic Python programming (e.g., ability to read and write simple Python scripts; understanding of introductory concepts like variables, loops, modules, conditionals, data types, and functions).
- A background in beginning level statistics would also be helpful to understand and deploy the ML models.
- Some experience with fundamental statistics and algebra, including an understanding of data distributions (e.g., normal distribution) measures of central tendency and variability (e.g., mean and standard deviation) and basic linear equations.
- Udacity also recommends basic familiarity with fundamental machine learning concepts (such as feature engineering and supervised vs. unsupervised learning) and classic machine learning algorithms (such as linear regression and k-means clustering).
- An understanding of the basics of Azure and Docker/Container experience
- If you'd like to prepare for this Nanodegree program, check out our [Introduction to Machine Learning](#) and [AI Programming with Python courses](#).


IN COLLABORATION WITH



Estimated Time:
3 Months at
5-10 hours / week



Prerequisites:
Prior experience
with Python,
Machine Learning,
and Statistics



Flexible Learning:
Self-paced, so
you can learn on
the schedule that
works best for you.



Need Help?
udacity.com/advisor
Discuss this program
with an enrollment
advisor.

Course 1: Using Azure Machine Learning

Machine learning is a critical business operation for many organizations. Learn how to configure machine learning pipelines in Azure. Identify use cases for Automated Machine Learning. Use the Azure ML SDK to design, create, and manage machine learning pipelines in Azure.

Project 1 Optimizing an ML Pipeline in Azure

Throughout the course, we cover many different ways to work with data and machine learning. It can be quite challenging to decide what method to use - building your own machine learning pipeline, leveraging AutoML, hyperparameter tuning, and so on. In this project, students use scikit-learn, Hyperdrive, and AutoML to understand the costs and benefits of each methodology. First, students will construct a pipeline from scikit-learn, using the Azure ML SDK to import data from a URL. Then, students will configure a Hyperdrive run for their scikit-learn pipeline to find the optimal hyperparameters. Students will then use the same dataset for an AutoML run to find an optimal model and set of hyperparameters. Finally, students write a README documenting their findings and comparing the differences, costs, and benefits of the different methods they've used.

LEARNING OUTCOMES

LESSON ONE

Introduction to Azure ML

- Understand why you should do ML in the cloud
- Understand when you should do ML in the cloud
- Analyze the customers of ML

LESSON TWO

Workspaces and the Azure ML Studio

- Interpret the Azure ML Platform
- Explain how to manage and choose compute resources
- Summarize the key components of Workspaces and Notebooks

LESSON THREE

Datastores and Datasets

- Analyze how to manage data
- Construct datasets
- Compose solutions to manage data drift and deal with sensitive data

LESSON FOUR

Training Models in Azure ML

- Experiment with the Designer
- Develop and manage pipelines
- Organize and run hyperparameter experiments

LESSON FIVE

The AzureML SDK

- Utilize data with the SDK
- Create pipelines
- Organize experiments

LESSON SIX

AutoML and Hyperparameter

- Design solutions with AutoML and the SDK
- Analyze model interpretation experiments
- Create portable ML models with ONNX

Course 2: Machine Learning Operations

Operationalizing Machine Learning is a set of best practices that are mostly inherited by the DevOps movement. In the past few years, it has become clear that shipping models into production in a reliable, reproducible, and automated way with a constant feedback loop is crucial. This is where all the DevOps principles come into play and is exactly what this course covers in detail.

Project 2
Operationalizing
Machine Learning

MLOps and its core features have been covered in this course in detail. This project will apply all the principles from the lessons to get a model trained with AutoML and deployed into a production environment.

This project covers a lot of the key concepts of operationalizing Machine Learning, from selecting the appropriate targets for deploying models, to enabling Application Insights, identifying problems in logs, and harnessing the power of Azure’s Pipelines. All these concepts are part of core DevOps pillars that will allow you to demonstrate solid skills for shipping machine learning models into production.

LEARNING OUTCOMES

LESSON ONE

Enabling Security

- Create a Service Principal account for different types of roles
- Determine what the differences are in various forms of authentication
- Use a specific type of authentication when selecting deployment settings

LESSON TWO

Deploy a ML model

- Use a production environment for deployment
- Enable authentication in the deployment cluster
- Discover the differences between container-based deployment and kubernetes.

LEARNING OUTCOMES

LESSON ONE

ML Endpoints

- Use a proven tool to find what a baseline for performance is
- Gather information about an endpoint input to interact with it
- Find what potential issues can happen with incorrect input

LESSON TWO

Pipeline Automation

- Create a pipeline to further automation when training models
- Enable a REST API for the pipeline, so other services can interact with it

Capstone Project

The program capstone gives you the opportunity to use the knowledge you have obtained from this Nanodegree program to solve an interesting problem. You will have to use Azure's Automated ML and HyperDrive to solve a task. Finally, you will have to deploy the model as a webservice and test the model endpoint.

CapstoneProject

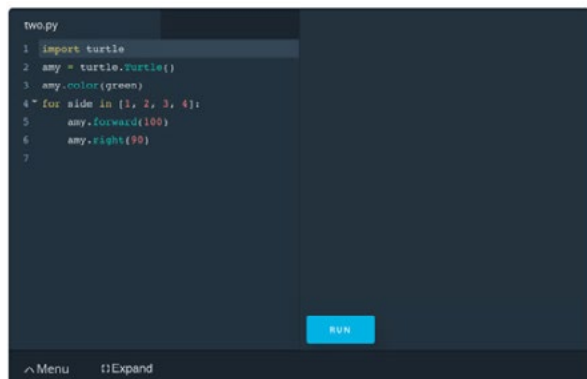


Training and Deploying a Machine Learning Model in Azure

You will be using both the hyperdrive and automl API from azureml to build this project. You can choose the model you want to train, and the data you want to use. However, the data you use needs to be external and not available in Azure's ecosystem. After you have chosen a dataset, you will have to import the dataset into your workspace. Subsequently, you will train a model on that dataset using automated ML and then train a custom model whose hyperparameters you have tuned using HyperDrive. The type of model you use is not important. You can use ML models through Scikit-learn or Deep Learning models like ANNs and CNNs through Keras, TensorFlow, or PyTorch for this part of the project.

After you have trained both the models, compare their performance, deploy the best model as a webservice and test the model endpoint.

This project will demonstrate your ability to use an external dataset in your workspace, train a model using the different tools available in the AzureML framework as well as your ability to deploy the model as a web service.

Our Classroom Experience



REAL-WORLD PROJECTS

Build your skills through industry-relevant projects. Get personalized feedback from our network of 900+ project reviewers. Our simple interface makes it easy to submit your projects as often as you need and receive unlimited feedback on your work.

KNOWLEDGE

Find answers to your questions with Knowledge, our proprietary wiki. Search questions asked by other students and discover in real-time how to solve the challenges that you encounter.

STUDENT HUB

Leverage the power of community through a simple, yet powerful chat interface built within the classroom. Use Student Hub to connect with your technical mentor and fellow students in your Nanodegree program.

WORKSPACES

See your code in action. Check the output and quality of your code by running them on workspaces that are a part of our classroom.

QUIZZES

Check your understanding of concepts learned in the program by answering simple and auto-graded quizzes. Easily go back to the lessons to brush up on concepts anytime you get an answer wrong.

CUSTOM STUDY PLANS

Work with a mentor to create a custom study plan to suit your personal needs. Use this plan to keep track of your progress toward your goal.

PROGRESS TRACKER

Stay on track to complete your Nanodegree program with useful milestone reminders.

Learn with the Best



Noah Gift

FOUNDER OF PRAGMATICAI LABS

Noah Gift teaches and consults at top universities and companies globally, including Duke and Northwestern. His areas of expertise are machine learning, MLOps, A.I., Data Science, Machine Learning, and Cloud Architecture. Noah has authored several bestselling books, including Python for DevOps.



Alfredo Deza

INSTRUCTOR

Alfredo Deza is a passionate software engineer, avid open source developer, Vim plugin author, photographer, and former Olympic athlete. He has rebuilt company infrastructure, designed shared storage, and replaced complex build systems, always in search of efficient and resilient environments.



Erick Galinkin

PRINCIPAL AI RESEARCHER | RAPID7

Erick Galinkin is a hacker and scientist specializing in Applying Artificial Intelligence to Cybersecurity. He also conducts academic research on machine learning theory and the interplay between algorithmic game theory and information security.



Soham Chatterjee

GRADUATE STUDENT AT THE NANYANG TECHNOLOGICAL UNIVERSITY

Soham is an Intel® Software Innovator and a former Deep Learning Researcher at Saama Technologies. He is currently a Masters by Research student at NTU, Singapore. His research is on Edge Computing, IoT and Neuromorphic Hardware.

All Our Nanodegree Programs Include:



EXPERIENCED PROJECT REVIEWERS

REVIEWER SERVICES

- Personalized feedback & line by line code reviews
- 1600+ Reviewers with a 4.85/5 average rating
- 3 hour average project review turnaround time
- Unlimited submissions and feedback loops
- Practical tips and industry best practices
- Additional suggested resources to improve



TECHNICAL MENTOR SUPPORT

MENTORSHIP SERVICES

- Questions answered quickly by our team of technical mentors
- 1000+ Mentors with a 4.7/5 average rating
- Support for all your technical questions



PERSONAL CAREER SERVICES

CAREER SUPPORT

- Resume support
- Github portfolio review
- LinkedIn profile optimization

Frequently Asked Questions

PROGRAM OVERVIEW

WHY SHOULD I ENROLL?

Businesses everywhere are mobilizing to create and implement the AI strategies that will transform industries in coming years, and they need engineers to do it. Data from LinkedIn indicates that AI specialists are among the most sought after roles that companies are looking for, with a 74 percent annual growth rate in hiring over the last four years. To stay in-demand at companies on the cutting edge of technology, engineers should prioritize developing their machine learning skill set.

The Machine Learning Engineer for Microsoft Azure Nanodegree Program, built in collaboration with Microsoft, offers you the chance to build the practitioner-level skills that companies across industries need. In the program, you'll strengthen your machine learning skills by training, validating, and evaluating machine learning models on the Azure platform, and complete a series of three real-world projects to add to your portfolio.

WHAT JOBS WILL THIS PROGRAM PREPARE ME FOR?

Students in the program will learn about machine learning algorithms and crucial deployment techniques, and will be equipped to fill roles at companies seeking machine learning engineers and AI specialists. These skills can also be applied in roles at companies that are looking for data scientists to introduce machine learning techniques into their organization.

HOW DO I KNOW IF THIS PROGRAM IS RIGHT FOR ME?

The Machine Learning Engineer for Microsoft Azure Nanodegree program is geared towards people who are interested in building and deploying a machine learning product or application. The program is a good fit for...

- Data scientists who are trying to expand their knowledge and application of ML techniques
- Software Developers who want to add ML concepts and techniques into their toolkit, or use Microsoft Azure for ML model development
- Other professionals who understand ML foundations but want to deepen their knowledge of and experience with practical applications of ML skills

ENROLLMENT AND ADMISSION

DO I NEED TO APPLY? WHAT ARE THE ADMISSION CRITERIA?

There is no application. This Nanodegree program accepts everyone, regardless of experience and specific background.

WHAT ARE THE PREREQUISITES FOR ENROLLMENT?

- Experience with basic Python programming (e.g., ability to read and write simple Python scripts; understanding of introductory concepts like variables, loops, modules, conditionals, data types, and functions).
- Some experience with fundamental statistics and algebra, including an understanding of data distributions (e.g., normal distribution) measures of central tendency and variability (e.g., mean and standard deviation) and basic linear equations.



FAQs Continued

- Udacity also recommends basic familiarity with fundamental machine learning concepts (such as feature engineering and supervised vs. unsupervised learning) and classic machine learning algorithms (such as linear regression and k-means clustering).
- An understanding of the basics of Azure and Docker/Container experience.
- If you'd like to prepare for this Nanodegree program, check out our [Introduction to Machine Learning](#) and [AI Programming with Python](#) courses.

IF I DO NOT MEET THE REQUIREMENTS TO ENROLL, WHAT SHOULD I DO?

To prepare, we recommend the [Introduction to Machine Learning](#) and [AI Programming with Python](#) programs, to build your comfortability with ML concepts and using python in an AI context.

TUITION AND TERM OF PROGRAM

HOW IS THIS NANODEGREE PROGRAM STRUCTURED?

The Machine Learning Engineer for Microsoft Azure Nanodegree program is comprised of content and curriculum to support three (3) projects. We estimate that students can complete the program in three (3) months working 5-10 hours per week.

Each project will be reviewed by the Udacity reviewer network. Feedback will be provided and if you do not pass the project, you will be asked to resubmit the project until it passes.

HOW LONG IS THIS NANODEGREE PROGRAM?

Access to this Nanodegree program runs for the length of time specified in the payment card on the Nanodegree program overview page. If you do not graduate within that time period, you will continue learning with month to month payments. See the [Terms of Use](#) for other policies around the terms of access to our Nanodegree programs.

SOFTWARE AND HARDWARE

WHAT SOFTWARE AND VERSIONS WILL I NEED IN THIS PROGRAM?

There are no specific hardware or software requirements for this program, other than those outlined on Udacity's general [Technology Requirements page](#).

