



THE SCHOOL OF ARTIFICIAL INTELLIGENCE

AWS Machine Learning Engineer



NANODEGREE SYLLABUS

Overview

AWS Machine Learning Engineer Nanodegree Program

IN COLLABORATION WITH



The goal of the AWS Machine Learning Engineer (MLE) Nanodegree program is to equip software developers/data scientists with the data science and machine learning skills required to build and deploy machine learning models in production using Amazon SageMaker. This program will focus on the latest best practices and capabilities that are enabled by Amazon SageMaker, including new model design/deployment features and case studies in which they can be applied to.

Educational Objectives

A graduate of this program will be able to:

- Create machine learning models in Sagemaker on datasets cleaned using AWS tools
- Deploy machine learning models to an API endpoint and integrate it into a full workflow
- Solve computer vision and natural language problems using fine-tuned deep neural networks
- Operationalize a machine learning pipeline using SageMaker to allow for training and deployment on industry-scale problems
- Select a machine learning challenge and propose a possible solution

Program Information



TIME

5 months
Study 5-10 hours/week



LEVEL

Intermediate



PREREQUISITES

- At least 40 hours of programming experience
- Familiarity with data structures like dictionaries and lists
- Experience with libraries like NumPy and pandas
- Knowledge of functions, variables, loops, and classes
- Exposure to Python through Jupyter Notebooks is recommended
- Experience with constructing and calling HTTP API endpoints is recommended

Basic knowledge of machine learning algorithms, including:

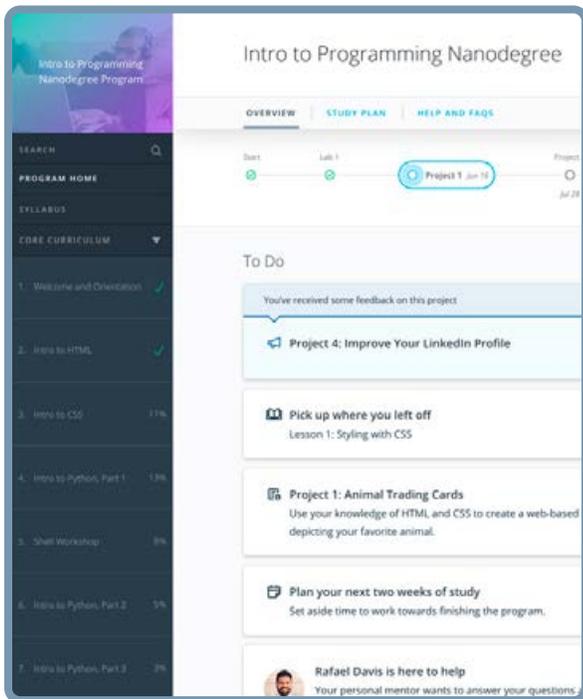
- Basic understanding of the machine learning workflow
- Basic theoretical understanding of ML algorithms such as linear regression, logistic regression, and neural network
- Basic understanding of model training and testing processes
- Basic knowledge of commonly used metrics for ML models evaluation such as accuracy, precision, recall, and mean square error (MSE)



LEARN MORE ABOUT THIS NANODEGREE

Contact us at
enterpriseNDs@udacity.com.

Our Classroom Experience



REAL-WORLD PROJECTS

Learners build new skills through industry-relevant projects and receive personalized feedback from our network of 900+ project reviewers. Our simple user interface makes it easy to submit projects as often as needed and receive unlimited feedback.

KNOWLEDGE

Answers to most questions can be found with Knowledge, our proprietary wiki. Learners can search questions asked by others and discover in real time how to solve challenges.

LEARNER HUB

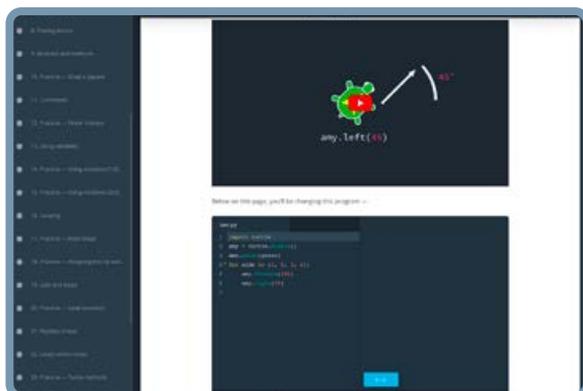
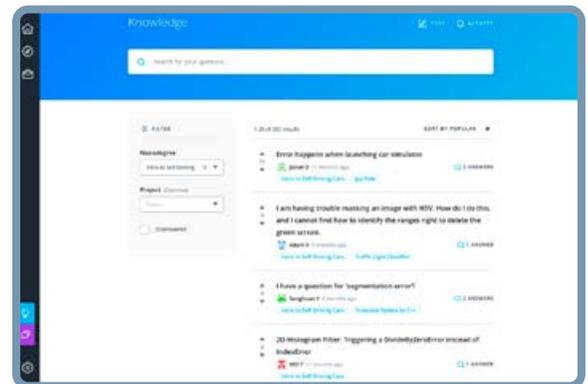
Learners leverage the power of community through a simple, yet powerful chat interface built within the classroom. Learner Hub connects learners with their technical mentor and fellow learners.

WORKSPACES

Learners can check the output and quality of their code by testing it on interactive workspaces that are integrated into the classroom.

QUIZZES

Understanding concepts learned during lessons is made simple with auto-graded quizzes. Learners can easily go back and brush up on concepts at any time during the course.



CUSTOM STUDY PLANS

Mentors create a custom study plan tailored to learners' needs. This plan keeps track of progress toward learner goals.

PROGRESS TRACKER

Personalized milestone reminders help learners stay on track and focused as they work to complete their Nanodegree program.

Learn with the Best



Matt Maybeno

PRINCIPAL SOFTWARE ENGINEER
DATA SCIENCE AND MACHINE
LEARNING

Matt Maybeno is a principal software engineer at SOCi. With a masters in bioinformatics from SDSU, he utilizes his cross domain expertise to build solutions in NLP and predictive analytics.



Joseph Nicolls

SENIOR MACHINE LEARNING
ENGINEER, BLUE HEXAGON

Joseph Nicolls is a senior machine learning scientist at Blue Hexagon. With a major in biomedical computation from Stanford University, he currently utilizes machine learning to build malware-detecting solutions at Blue Hexagon.



Charles Landau

TECHNICAL LEAD,
AI/ML GUIDEHOUSE

Charles Landau is a developer at Guidehouse, a management consulting company. Charles holds a MPA from George Washington University, where he focused on econometrics and regulatory policy, and holds a BA from Boston University. At Guidehouse, he supports data scientists and developers working on internal and client-facing ML platforms.



Soham Chatterjee

MULTI CLOUD ENGINEER

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.



Bradford Tuckfield

INDEPENDENT CONSULTANT

Bradford does independent consulting for machine learning projects related to manufacturing, law, pharmaceutical operations, and other fields. He also writes technical books about programming, algorithms, and data science.

Nanodegree Program Overview

Course 1: Introduction to Machine Learning

In this course, you'll start learning what machine learning is by being introduced to the high level concepts through AWS SageMaker. You'll begin by using SageMaker Studio to perform exploratory data analysis. Know how and when to apply the basic concepts of machine learning to real world scenarios. Create machine learning workflows, starting with data cleaning and feature engineering, to evaluation and hyperparameter tuning. Finally, you'll build new ML workflows with highly sophisticated models such as XGBoost and AutoGluon.

Project

Predict Bike Sharing Demand with AutoGluon

In this project, students will apply the knowledge and methods they learned in the Introduction to Machine Learning course to compete in a Kaggle competition. Using the AutoGluon framework, students will first train a baseline model, then improve their model through feature engineering and hyperparameter tuning. Finally, they'll submit their optimized model for a public Kaggle rank and write a report on their findings to showcase their work.



Nanodegree Program Overview



LESSON TITLE

LEARNING OUTCOMES

EXPLORATORY DATA ANALYSIS

- Use AWS SageMaker Studio to access datasets from S3 and perform data analysis functions using AWS tools
- Perform data analysis and feature engineering with Data Wrangler
- Perform data analysis and feature engineering with Pandas in SageMaker Studio
- Label new data for a dataset with Sagemaker ground truth

MACHINE LEARNING CONCEPTS

- Design a domain, model, and data outline for a case study
- Build a ML lifecycle and apply it to a dataset
- Differentiate between supervised and unsupervised models and apply them to an appropriate dataset
- Differentiate between regression and classification methods and apply them to an appropriate dataset

MODEL DEPLOYMENT WORKFLOW

- Load new dataset, create 3 data set types, and identify features/values in SageMaker
- Clean or create new features from a dataset
- Train (fit) a regression/classification model using scikit learn
- Evaluate a trained model using methods like mse, rmse, r2, accuracy, f1, and precision
- Tune a model's hyper parameters to achieve a better result

ALGORITHMS AND TOOLS

- Train, test, and optimize a linear model, tree-based model, XGBoost model, and AutoGluon Tabular prediction model
- Create a model using Sagemaker Jumpstart

Nanodegree Program Overview

Course 2: Developing Your First ML Workflow

In order to execute on machine learning's versatile capabilities, we need to have the infrastructure to execute our ML operations. With the easy availability of managed infrastructure from AWS, we can dynamically create the necessary resources to train, deploy, and evaluate our models. In this course you will learn how to create general machine learning workflows on AWS.

You'll begin with an introduction to the general principles of machine learning engineering. From there, you'll learn the fundamentals of SageMaker to train, deploy, and evaluate a model. Following that, you'll learn how to create a machine learning workflow on AWS utilizing tools like Lambda and Step Functions. Finally, you'll learn how to monitor machine learning workflows with services like Model Monitor and Feature Store. With all this, you'll have all the information you need to create an end-to-end machine learning pipeline.

Project

Build a ML Workflow on SageMaker

In this project, students will develop an end-to-end ML Workflow on SageMaker, Lambda, and Step Functions. Students will showcase their model deployment capabilities with SageMaker Model Endpoints and Lambda, and their workflow monitoring capabilities with SageMaker Model Monitor and Step Functions. At the end of the project, students will be able to demonstrate building a scalable ML workflow on SageMaker.

Nanodegree Program Overview



LESSON TITLE

LEARNING OUTCOMES

INTRODUCTION TO MLE

- Understand the prerequisites
- Describe key business stakeholders
- Understand the history of MLE
- Describe when to use MLE

SAGEMAKER ESSENTIALS

- Launch training jobs within SageMaker
- Deploy an endpoint that can perform inference on live data
- Evaluate datasets with batch transform jobs
- Perform custom processing jobs on raw data

DESIGNING YOUR OWN WORKFLOW

- Create Lambda functions
- Trigger Lambda functions utilizing both the SDK and other AWS Services
- Design and execute a workflow utilizing State Machines
- Learn about the use cases for SageMaker Pipelines

MONITORING A ML WORKFLOW

- Use SageMaker Feature Store to serve and monitor model data
- Configure SageMaker Model Monitor to generate and track metrics about our models
- Use Clarify to explain model predictions and surface biases in models

Nanodegree Program Overview

Course 3: Deep Learning Topics within Computer Vision and NLP

As more machine learning products are being deployed, machine learning engineering is becoming a very important and sought after skill in the industry. Building infrastructures for training, deployment, and monitoring of deep learning models is different from building other software systems. In this course you will learn how to train, finetune and deploy deep learning models using Amazon SageMaker.

You'll begin by learning what deep learning is, where it is used and the tools used by deep learning engineers. Next we will learn about artificial neurons and neural networks and how to train them. After that we will learn about advanced neural network architectures like convolutional neural networks and BERT as well as how to finetune them for specific tasks. Finally, you will learn about Amazon SageMaker and you will take everything you learned and do them in SageMaker Studio.

Project

Image Classification using AWS SageMaker

In this project, students will be using AWS Sagemaker to finetune a pretrained model that can perform image classification. Students will have to use Sagemaker profiling, debugger, hyperparameter tuning and other good ML engineering practices to finish this project. To finish this project, students will have to perform tasks and use tools that a typical ML Engineer does as a part of their job.

Nanodegree Program Overview



LESSON TITLE

LEARNING OUTCOMES

INTRODUCTION TO DEEP LEARNING TOPICS WITHIN COMPUTER VISION AND NLP

- Understand the need and importance of deep learning
- Learn the history of deep learning and the business stakeholders in a deep learning project
- Learn the tools used by deep learning engineers

INTRODUCTION TO DEEP LEARNING

- Understand the workings of artificial neurons and neural networks
- Understand how to set cost functions and optimizers to train neural networks
- Build and train a neural network on an image classification task

COMMON MODEL ARCHITECTURE TYPES AND FINE-TUNING

- Understand how advanced neural network architectures like convolutional neural networks and transformer based models work
- Finetune a pretrained model on a different task
- Understand the important of hyperparameter tuning for training (and fine-tuning) deep neural networks

DEPLOY DEEP LEARNING MODELS ON SAGEMAKER

- Finetune models for image and text classification using SageMaker JumpStart
- Debug and profile training jobs using SageMaker Debugger
- Tune hyperparameters when training a model
- Package a model in a Dockerfile for deployment

Nanodegree Program Overview

Course 4: Operationalizing Machine Learning Projects on SageMaker

This course covers advanced topics related to deploying professional machine learning projects on SageMaker. It also covers security applications. You will learn how to maximize output while decreasing costs. You will also learn how to deploy projects that can handle high traffic, and how to work with especially large datasets.

Project

Operationalizing an AWS ML Project

In this project, students will start with a machine learning project that accomplishes computer vision tasks. Students will deploy the project on AWS and add several important features: cost minimization, security, and redeployment on a separate server. This project will prepare students to successfully deploy professional projects in industrial applications.



Nanodegree Program Overview



LESSON TITLE

LEARNING OUTCOMES

MANAGE COMPUTE RESOURCES IN AWS ACCOUNTS TO ENSURE EFFICIENT UTILIZATION

- Keep costs low in AWS machine learning projects
- Use spot instances for efficiency
- Turn off resources when they're not being used
- Check costs to ensure they remain low

TRAIN MODELS ON LARGE-SCALE DATASETS USING DISTRIBUTED TRAINING

- Perform multi-instance training
- Use distributed data to improve performance
- Create and interpret manifest files
- Choose the best data stores for projects

CONSTRUCT PIPELINES FOR HIGH-THROUGHPUT, LOW-LATENCY MODELS

- Set up Lambda functions for AWS projects
- Configure endpoints for auto-scaling
- Set up concurrency for Lambda functions
- Create feature stores for data imports

DESIGN SECURE MACHINE LEARNING PROJECTS IN AWS

- Resolve security issues using IAM settings
- Set up a virtual private cloud for security
- Manage security in SageMaker

Nanodegree Program Overview

CAPSTONE PROJECT: Inventory Monitoring at Distribution Centers

Distribution centers often use robots to move objects as a part of their operations. Objects are carried in bins where each bin can contain multiple objects. In this project, students will have to build a model that can count the number of objects in each bin. A system like this can be used to track inventory and make sure that delivery consignments have the correct number of items.

To build this project, students will have to use AWS Sagemaker and good machine learning engineering practices to fetch data from a database, preprocess it and then train a machine learning model. This project will serve as a demonstration of end-to-end machine learning engineering skills that will be an important piece of their job-ready portfolio.



Our Nanodegree Programs Include:



Pre-Assessments

Our in-depth workforce assessments identify your team's current level of knowledge in key areas. Results are used to generate custom learning paths designed to equip your workforce with the most applicable skill sets.



Dashboard & Progress Reports

Our interactive dashboard (enterprise management console) allows administrators to manage employee onboarding, track course progress, perform bulk enrollments and more.



Industry Validation & Reviews

Learners' progress and subject knowledge is tested and validated by industry experts and leaders from our advisory board. These in-depth reviews ensure your teams have achieved competency.



Real World Hands-on Projects

Through a series of rigorous, real-world projects, your employees learn and apply new techniques, analyze results, and produce actionable insights. Project portfolios demonstrate learners' growing proficiency and subject mastery.

Our Review Process

Real-life Reviewers for Real-life Projects

Real-world projects are at the core of our Nanodegree programs because hands-on learning is the best way to master a new skill. Receiving relevant feedback from an industry expert is a critical part of that learning process, and infinitely more useful than that from peers or automated grading systems. Udacity has a network of over 900 experienced project reviewers who provide personalized and timely feedback to help all learners succeed.



Vaibhav
UDACITY LEARNER

"I never felt overwhelmed while pursuing the Nanodegree program due to the valuable support of the reviewers, and now I am more confident in converting my ideas to reality."

now at
CODING VISIONS INFOTECH

All Learners Benefit From:



Line-by-line feedback for coding projects



Industry tips and best practices



Advice on additional resources to research



Unlimited submissions and feedback loops

How it Works

Real-world projects are integrated within the classroom experience, making for a seamless review process flow.

- Go through the lessons and work on the projects that follow
- Get help from your technical mentor, if needed
- Submit your project work
- Receive personalized feedback from the reviewer
- If the submission is not satisfactory, resubmit your project
- Continue submitting and receiving feedback from the reviewer until you successfully complete your project

About our Project Reviewers

Our expert project reviewers are evaluated against the highest standards and graded based on learners' progress. Here's how they measure up to ensure your success.

900+

Expert Project Reviewers

Are hand-picked to provide detailed feedback on your project submissions.

1.8M

Projects Reviewed

Our reviewers have extensive experience in guiding learners through their course projects.

3

Hours Average Turnaround

You can resubmit your project on the same day for additional feedback.

4.85 /5

Average Reviewer Rating

Our learners love the quality of the feedback they receive from our experienced reviewers.



UDACITY
FOR ENTERPRISE

Udacity © 2021

2440 W El Camino Real, #101
Mountain View, CA 94040, USA - HQ

For more information visit: www.udacity.com/enterprise