UDACITY
FOR ENTERPRISE

THE SCHOOL OF ARTIFICIAL INTELLIGENCE

Intel® Edge AI for IoT Developers

NANODEGREE SYLLABUS

Overview

This Nanodegree Program is built in partnership with intel

Edge AI applications are revolutionizing the IoT industry by bringing fast, intelligent behavior to the locations where it is needed. In this Nanodegree program, you will learn how to develop and optimize Edge AI systems, using the Intel® Distribution of OpenVINO™ Toolkit. A graduate of this program will be able to:

- Leverage the Intel® Distribution of OpenVINO™ Toolkit to fast-track development of high-performance computer vision and deep learning inference applications.
- Run pre-trained deep learning models for computer vision on-prem.
- Identify key hardware specifications of various hardware types (CPU, VPU, FPGA, and Integrated GPU).
- Utilize Intel® DevCloud for the Edge to test model performance on various hardware types (CPU, VPU, FPGA, and Integrated GPU).

This program consists of 3 courses and 3 projects. Each project you build will be an opportunity to demonstrate what you've learned in the course, and will demonstrate to potential employers that you have skills in these areas.This program is comprised of 4 courses and 4 projects. Each project you build will be an opportunity to demonstrate what you've learned in the lesson, and will demonstrate to potential employers that you have skills in these areas.

Program Information

TIME
3 months
Study 10 hours/week

LEVEL
Specialist

PREREQUISITES
- Intermediate Python
- Deep Learning
- Command Line
- OpenCV

HARDWARE/SOFTWARE REQUIRED
Computer running a 64-bit operating system that has 6th or newer generation of Intel® processor running either Win, Ubuntu or (copy from Intel® Distribution of OpenVINO™ Toolkit).

LEARN MORE ABOUT THIS NANODEGREE
Contact us at enterpriseNDs@udacity.com.

Our Classroom Experience

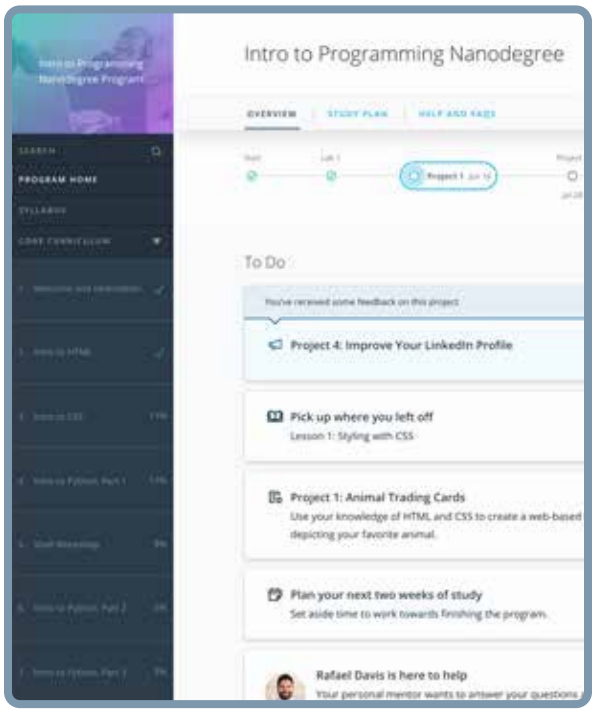

REAL-WORLD PROJECTS

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 anytime during the course.

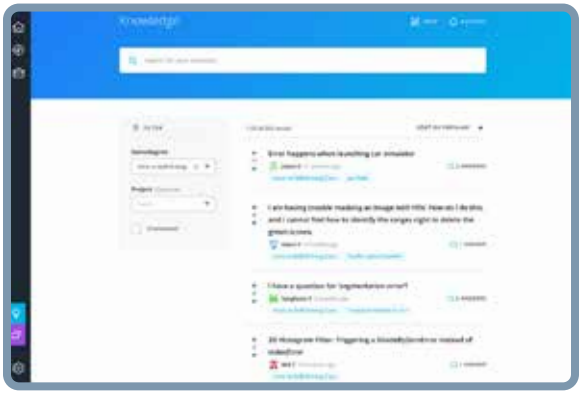

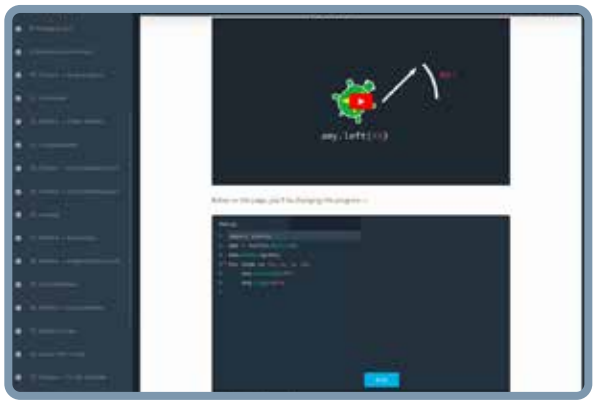

CUSTOM STUDY PLANS

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

Stewart Christie

COMMUNITY MANAGER, IOT DEVELOPER PROGRAM AT INTEL

Stewart has been with Intel for the past 20+ years. He is an experienced Technical Marketer with extensive hardware and software knowledge. He has comprehensive understanding of designing embedded electronic hardware and embedded firmware for various applications.

Michael Virgo

SENIOR CURRICULUM MANAGER AT UDACITY

After beginning his career in business, Michael utilized Udacity Nanodegree programs to build his technical skills, eventually becoming a Self-Driving Car Engineer at Udacity before switching roles to work on curriculum development for a variety of AI and Autonomous Systems programs.

Soham Chatterjee

GRADUATE STUDENT AT NTU

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

Vaidheeswaran Archana

GRADUATE STUDENT AT NUS

Archana is a graduate student at the National University of Singapore. She is currently pursuing her research in Deep Learning and Smart Grids, under Professor Dipti Srinivasan. Archana is an Intel Software Innovator and a former Deep Learning Engineer at Saama Technologies.

Course 1: Edge AI Fundamentals with OpenVINO™

Leverage a pre-trained model for computer vision inferencing. You will convert pre-trained models into the framework agnostic intermediate representation with the Model Optimizer, and perform efficient inference on deep learning models through the hardware-agnostic Inference Engine. Finally, you will deploy an app on the edge, including sending information through MQTT, and analyze model performance and use cases.

Project: Deploy a People Counter App at the Edge

In this project, you will utilize the Intel® Distribution of the OpenVINO™ Toolkit to build a People Counter app for inference at the edge. You will investigate different pre-trained models for person detection, and then convert the best model for optimized inference. The model will be deployed on the edge, such that only data on 1) the number of people in the frame, 2) time those people spent in frame, and 3) the total number of people counted are sent to a web server; inference will be done on the local machine.

You will need to develop a method to compare the performance of their models before and after use of the OpenVINO toolkit for optimization for edge deployment. You will also examine potential use cases for their deployed people counter app.

LESSON TITLE	LEARNING OUTCOMES
LEVERAGING PRE-TRAINED MODELS	• Leverage a pre-trained model for computer vision inferencing.
THE MODEL OPTIMIZER	• Convert pre-trained models into the framework agnostic intermediate representation with the Model Optimizer.
THE INFERENCE ENGINE	• Perform efficient inference on deep learning models through the hardware-agnostic Inference Engine.
DEPLOYING AN EDGE APP	• Deploy an app on the edge, including sending information through MQTT, and analyze model performance and use cases.

Course 2: Hardware for Computer Vision Deep Learning Application Deployment

Grow your expertise in choosing the right hardware. Identify key hardware specifications of various hardware types (CPU, VPU, FPGA, and Integrated GPU). Utilize the DevCloud to test model performance and deploy power-efficient deep neural network inference on on the various hardware types. Finally, you will distribute workload on available compute devices in order to improve model performance.

Project: Design a Smart Queuing System

In this project, you will be given a real-world scenario of building a queuing system for three different clients in three different industry sectors. The sectors will consist of retail, manufacturing, and transportation. Each client will have their own set of constraints and requirements. You'll use your knowledge of hardware specifications to identify which hardware types might work, and then you'll test the application using the Intel® DevCloud to see which hardware performs best. Finally, after reviewing your test results and considering the constraints and requirements of the client, you will propose a hardware solution and justify your selection.

LESSON TITLE	LEARNING OUTCOMES
INTRODUCTION TO HARDWARE AT THE EDGE	• Describe the importance of selecting the right hardware and the process involved in doing so.
CPU AND INTEGRATED GPU	• Identify the key specifications of Intel® CPUs and Integrated GPUs. • Use the Intel® Devcloud for the Edge for running deep learning models on the CPU and Integrated GPU.

Nanodegree Program Overview

LESSON TITLE	LEARNING OUTCOMES
VISION PROCESSING UNITS	• Identify the key specifications of Intel® VPUs. • Use the Intel® DevCloud for the Edge for running deep learning models on the VPU. • Use the MULTI Plugin to get more consistent performance.
FIELD PROGRAMMABLE GATE ARRAYS	• Identify the key specifications of Intel® FPGAs. • Use the Intel® DevCloud for the Edge for running deep learning models on the FPGA. • Use the HETERO Plugin to enable efficient hardware utilization.

Course 3: Optimization Techniques and Tools for Computer Vision Deep Learning Applications

Learn how to optimize your model and application code to reduce inference time when running your model at the edge. Use different software optimization techniques to improve the inference time of your model. Calculate how computationally expensive your model is. Use DL Workbench to optimize your model and benchmark the performance of your model. Use a VTune amplifier to find and fix hotspots in your application code. Finally, package your application code and data so that it can be easily deployed to multiple devices.

Project: Build a Computer Pointer Controller

In this project, you will use models available in the OpenVINO toolkit to control your computer pointer using your eye gaze. You will first have to identify faces and extract a face from an input video stream captured from a webcam or a video file. Then you will need to extract facial landmarks and also use a head pose estimation model to find the orientation of the extracted face. Using the head pose and facial landmarks, you will find the orientation of the eye gaze using a gaze estimation model. Finally, you will need to move the mouse pointer in the direction of the eye gaze. This project will demonstrate your ability to run multiple models in the same machine, and coordinate and optimize the flow of data between those models.

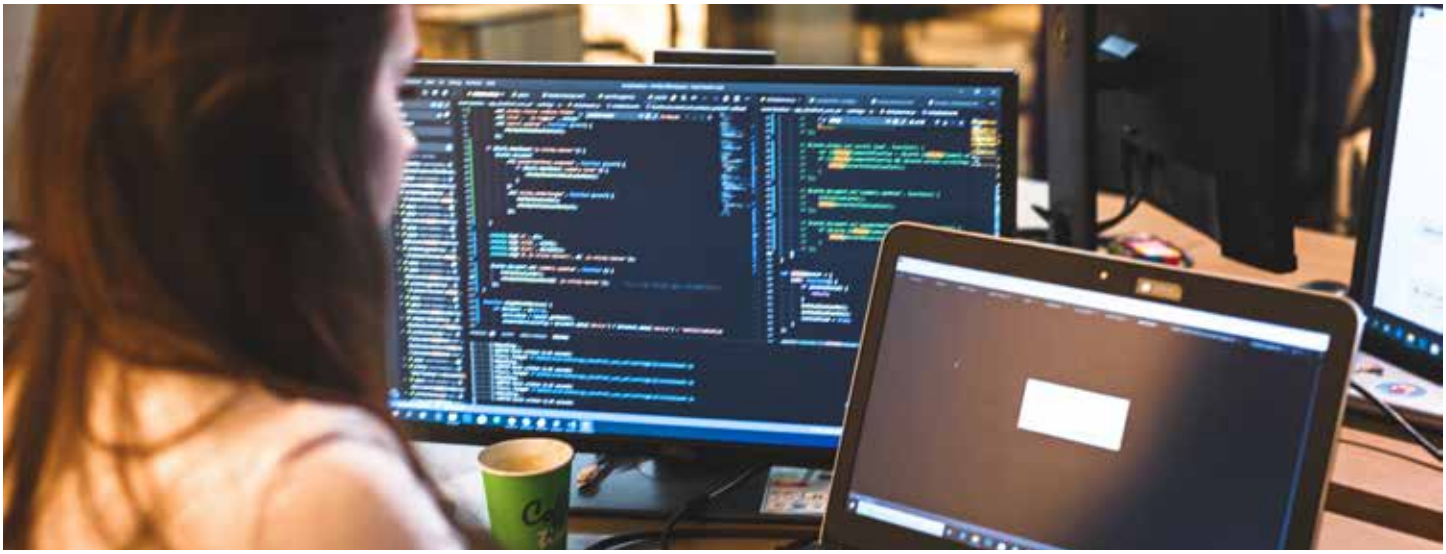

Nanodegree Program Overview

LESSON TITLE	LEARNING OUTCOMES
INTRODUCTION TO SOFTWARE OPTIMIZATION	• Describe why Software Optimization is important. • Identify the different fundamental optimization techniques. • Use different metrics to measure your model performance. • Identify when and how to use optimization techniques.
REDUCING MODEL OPERATIONS	• Calculate the number of operations in a model. • Implement optimization techniques that improve performance by reducing the number of model operations. • •Use OpenVINO to measure the effects of different optimization techniques on the performance of your model.
REDUCING MODEL SIZE	• Implement optimization techniques that improve performance by reducing the number of model operations. • Use DL Workbench to quantize and measure the performance of your model.
OTHER SOFTWARE OPTIMIZATION TECHNIQUES	• Use VTune Amplifier to measure hotspots in your application code. • Package your application code and data.

Our Nanodegree Programs Include:

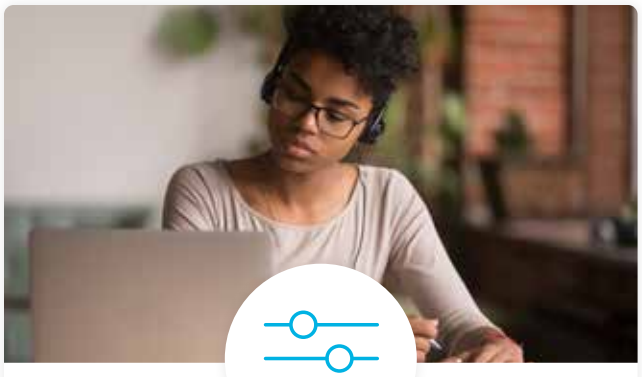

Pre-Assessments

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

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

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