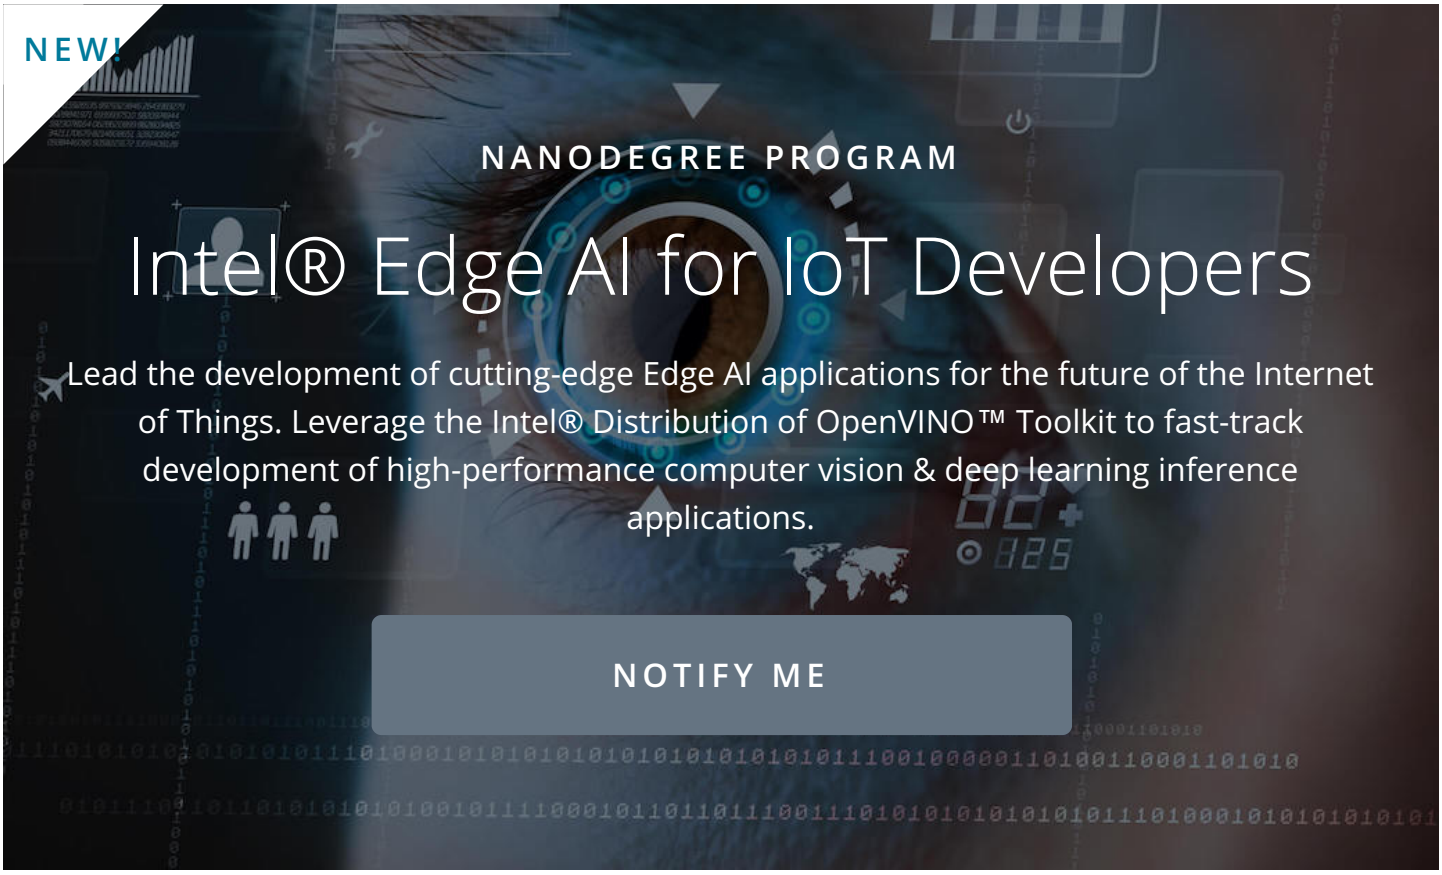

IN COLLABORATION WITH

intel®

What You Will Learn

DOWNLOAD SYLLABUS

Go to Course

SYLLABUS

Intel® Edge AI for IoT Developers

Lead the development of cutting-edge Edge AI applications that are the future of the Internet of Things. Leverage the Intel® Distribution of OpenVINO™ Toolkit to fast-track development of high-performance computer vision and deep learning inference applications.

Related Nanodegrees

— HIDE DETAILS

PREREQUISITE KNOWLEDGE

This program requires intermediate knowledge of Python, and experience with Deep Learning, Command Line, and OpenCV. See detailed requirements.

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

DEPLOY A PEOPLE COUNTER AT THE EDGE

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 Intel® DevCloud for the Edge 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.

DESIGN A SMART QUEUING SYSTEM

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 the 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.

BUILD A COMPUTER POINTER CONTROLLER

All Our Programs Include

Real-world projects from industry experts

With real world projects and immersive content built in partnership with top tier companies, you'll master the tech skills companies want.

Technical mentor support

Our knowledgeable mentors guide your learning and are focused on answering your questions, motivating you and keeping you on track.

Career services

You'll have access to resume support, Github portfolio review and LinkedIn profile optimization to help you advance your career and land a high-paying role.

Flexible learning program

Get a custom learning plan tailored to fit your busy life. Learn at your own pace and reach your personal goals on the schedule that works best for you.

Full List Of Offerings Included:

CLASS CONTENT	
Content co-created with Intel®	✓
Real-world projects	✓
Project reviews	✓
Project feedback from experienced reviewers	✓

STUDENT SERVICES	
Technical mentor support NEW	✓
Student community IMPROVED	✓
CAREER SERVICES	
Resume support	✓
Github review	✓
Linkedin profile optimization	✓

Succeed with Personalized Services

We provide services customized for your needs at every step of your learning journey to ensure your success!

Project Reviewers | Technical Mentor Support

Get timely feedback on your projects

REVIEWS BY THE NUMBERS

1,400+ project reviewers

2.7M projects reviewed

88/100 reviewer rating

1.1 hours avg project review turnaround time

REVIEWER SERVICES

- Personalized feedback
- Unlimited submissions and feedback loops
- Practical tips and industry best practices
- Additional suggested resources to improve

Learn with the Best

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

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

variety of AI and Autonomous Systems

Computing, IoT and Neuromorphic Hardware.

Related Nanodegree Programs

Data Structures and Algorithms

NANODEGREE PROGRAM

Get hands-on practice with over 100 data structures and algorithm exercises and technical mentor support when needed to help prepare you for interviews and on-the-job scenarios.

INTERMEDIATE

Intro to Machine Learning with TensorFlow

NANODEGREE PROGRAM

Build a solid foundation in Supervised, Unsupervised, and Deep Learning. Then, use these skills to test and deploy machine learning models in a production environment.

INTERMEDIATE

Program Details

PROGRAM UPDATE

Note: this program is not currently accepting new enrollments.

At Udacity it's important to us to present the most current technology possible. We have temporarily stopped enrollments for the Intel® Edge AI for IoT Developers Nanodegree Program as we are in the process of updating the program to be compatible with the latest release of DevCloud. You can register to be notified when the program is relaunched by clicking the "Notify Me" button above and submitting your information. We are unable to share a specific timeline for when the program will be available next, but hope to do so as soon as possible.

PROGRAM OVERVIEW - WHY SHOULD I TAKE THIS PROGRAM?

Why should I enroll?

70% of data being created is at the edge, and only half of that will go to the public cloud; the rest will be stored and processed at the edge, which requires a different kind of developer. Demand for professionals with the Edge AI skills will be immense, as the Edge Artificial Intelligence (AI) software market size is forecasted to grow from $355 Million in 2018, to $1.15 billion by 2023, at an Annual Growth Rate of 27%.(MarketsandMarkets) In the Edge AI for IoT Developers Nanodegree program, you'll leverage the potential of edge computing and use the Intel® Distribution of OpenVINO™ Toolkit to fast-track development of high-performance computer vision and deep learning inference applications.

Computer Vision is a fast-growing technology being deployed in nearly every industry from factory floors to amusement parks to shopping malls, smart buildings, and smart homes. It

is also driving the evolution of machine learning and human interactions with intelligent systems. Additional applications include drones, security cameras, robots, facial recognition on cell phones, self-driving vehicles, and more, which means these industries and more all need developers with computer vision and deep learning IoT experience.

What jobs will this program prepare me for?

This Nanodegree program will prepare you for roles such as IoT Developer, IoT Engineer, Deep Learning Engineer, Machine Learning Engineer, AI Specialist, VPU/CPU/FPGA Developer and more for companies and organizations looking to innovate their hardware on the Edge.

How do I know if this program is right for me?

If you are an enterprise developer and/or professional developer interested in advanced learning, specifically deep learning and computer vision, this program is right for you. Additionally if you have a background as an IoT Application Prototyper, IoT Application Implementer, IoT System Prototyper, or an IoT System Implementer, or in heterogeneous architectures as a Device Developer, Application Prototyper, Algorithm Developer, Solution Developer, or in security as an Architect/Planner, Security Specialist, or a Protocol Implementer, this program is a good fit.

What is Edge AI? What are some applications of this technology?

Edge Computing runs processes locally on the device itself, instead of running them in the cloud. This reduced computing time allows data to be processed much faster, removes the security risk of transferring the data to a cloud-based server, and reduces the cost of data transfer, as well as the risks of bandwidth outages disrupting performance.
Computer vision and AI at the edge are becoming instrumental in powering everything from factory assembly lines and retail inventory management, to hospital urgent care

medical imaging equipment like X-ray and CAT scans. Drones, security cameras, robots, facial recognition on cell phones, self-driving vehicles, and more all utilize this technology as well.

According to IEEE Innovation at Work, "By 2020, approximately 20+ billion devices will likely be connected via the Internet of Things (IoT), creating incredible amounts of data every minute. The time it takes to move data to the cloud, perform service on it and then move it back to devices is far too long to meet the increasing needs of the IoT. Unlike cloud computing, which relies on a single data center, edge computing works with a more distributed network, eliminating the round-trip journey to the cloud and offering real-time responsiveness and local authority. It keeps the heaviest traffic and processing closest to the end-user application and devices – smartphones, tablets, home security systems, and more – that generate and consume data. This dramatically reduces latency and leads to real-time, automated decision-making." (IEEE)

What is the Intel® DevCloud for the Edge?

The Intel® DevCloud for the Edge allows you to actively prototype and experiment with AI workloads for computer vision on Intel® hardware.

You have full access to hardware platforms hosted in our cloud environment, designed specifically for deep learning. You can test the performance of your models using the Intel® Distribution of OpenVINO™ Toolkit and combinations of CPUs, GPUs, VPUs such as the Intel® Neural Compute Stick 2 (NCS2) and FPGAs, such as the Intel® Arria® 10. The Intel® DevCloud for the Edge contains a series of Jupyter *notebook tutorials and examples preloaded with everything you needed to quickly get started.*

This includes trained models, sample data and executable code from the Intel® Distribution of OpenVINO™ Toolkit as well as other tools for deep learning. These notebooks are designed to help you quickly learn how to implement deep learning applications to enable compelling, high-performance solutions. Intel® has AI hardware waiting for your prototyping of edge inference jobs.

No hardware setup is required on your end. The Intel® DevCloud for the Edge utilizes Jupyter

Notebooks to execute code directly within the Web browser. Jupyter *is a browser-based development environment which allows you to run code and immediately visualize results. You can prototype innovative computer vision solutions in our cloud environment, then execute your code on any of Intel's® available combination of hardware resources.*

What is the Intel® Distribution of OpenVINO™ Toolkit and the Deep Learning Workbench?

You are able to deploy high performance, deep learning inference with the Intel® Distribution of OpenVINO™ Toolkit.

The Intel® Distribution of OpenVINO™ Toolkit allows you to harness the full potential of AI and computer vision across multiple Intel® architectures to enable new and enhanced use cases in health and life sciences, retail, industrial and more. Develop applications and solutions that emulate human vision with the Intel® Distribution of OpenVINO™ toolkit. Based on convolutional neural networks (CNN), the toolkit extends workloads across Intel® hardware (including accelerators) and maximizes performance.

- Enables deep learning inference from edge to cloud
- Accelerates AI workloads, including computer vision, audio, speech, language, and recommendation systems
- Supports heterogeneous execution across Intel® architecture and AI accelerators—CPU, iGPU, Intel® Movidius™ Vision Processing Unit (VPU), FPGA, and Intel® Gaussian & Neural Accelerator (Intel® GNA)—using a common API
- Speeds up time to market via a library of functions and pre-optimized kernels
- Includes optimized calls for OpenCV, OpenCL™ kernels, and other industry tools and libraries

The DL Workbench, part of the Intel® Distribution of OpenVINO™ Toolkit, is a web-based

graphical environment that enables users to visualize a simulation of performance of deep learning models and datasets on various Intel® architecture configurations (CPU, GPU, VPU). In addition, users can automatically fine-tune the performance of an Intel® Distribution of OpenVINO™ Toolkit model by reducing the precision of certain model layers (calibration) from FP32 to INT8. Additional tuning algorithms will be supported in future releases.

What makes the Intel® Edge AI for IoT Developers Nanodegree program unique?

The Intel® Distribution of OpenVINO™ Toolkit is for developers looking to deploy deep learning models on hardware with Intel® chips. Students will be able to interact with Intel's® IoT development platform to optimize the performance of their hardware using the DL Workbench. Through Udacity's interactive workspaces, you'll be able to send jobs to Intel® DevCloud for the Edge and see how different hardware performs in real time. Deploying AI models on the Edge requires a particular set of tools that providers such as Intel® have built. Through Udacity's hands-on exercises that integrate with Intel's® platform, students will be able to actually practice testing AI model performance on hardware without needing access to the hardware.
The Intel® DevCloud for the Edge is a cloud-based platform that lets you deploy machine learning models on hardware in the cloud before you purchase the actual hardware so you test and compare the performance of different hardware.

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?

To succeed in this program, students should have the following:

- Intermediate knowledge of programming in Python
- Experience with training and deploying deep learning models
- Familiarity with different DL layers and architectures (CNN based)
- Familiarity with the command line (bash terminal)
- Experience using OpenCV

If I do not meet the requirements to enroll, what should I do?

There are a few courses that can help prepare you for the program:

- For Python Experience:AI Programming with Python or Intro to Programming
- For C++ Experience (this will go in more depth than you will need): C++
- For AI Modeling: Intro to Machine Learning with Pytorch or Intro to Machine Learning with TensorFlow
- For Computer Vision Experience: Computer Vision

TUITION AND TERM OF PROGRAM

How is this Nanodegree program structured?

The Intel® Edge AI for IoT Developers 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 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?

You will have access to this Nanodegree program for as long as your subscription remains active. The estimated time to complete this program can be found on the webpage and in the syllabus, and is based on the average amount of time we project that it takes a student to complete the projects and coursework. See the Terms of Use and FAQs for other policies regarding the terms of access to our Nanodegree programs.

Can I switch my start date? Can I get a refund?

Please see the Udacity Program Terms of Use and FAQs for policies on enrollment in our programs.

I have graduated from the Intel® Edge AI for IoT Developers Nanodegree program, but I want to keep learning. Where should I go from here?

Many of our graduates continue on to our Machine Learning Engineer Nanodegree program, and after that, to the Self-Driving Car Engineer and Artificial Intelligence Nanodegree programs.

SOFTWARE AND HARDWARE - WHAT DO I NEED FOR THIS PROGRAM?

What software and versions will I need in this program?

You will need a 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).

You will also need to:

- Install Intel® Distribution of OpenVINO™ Toolkit (OpenVINO™ 2020.1) on your local environment.
- Install Intel's® Deep Learning Workbench

Go to Course