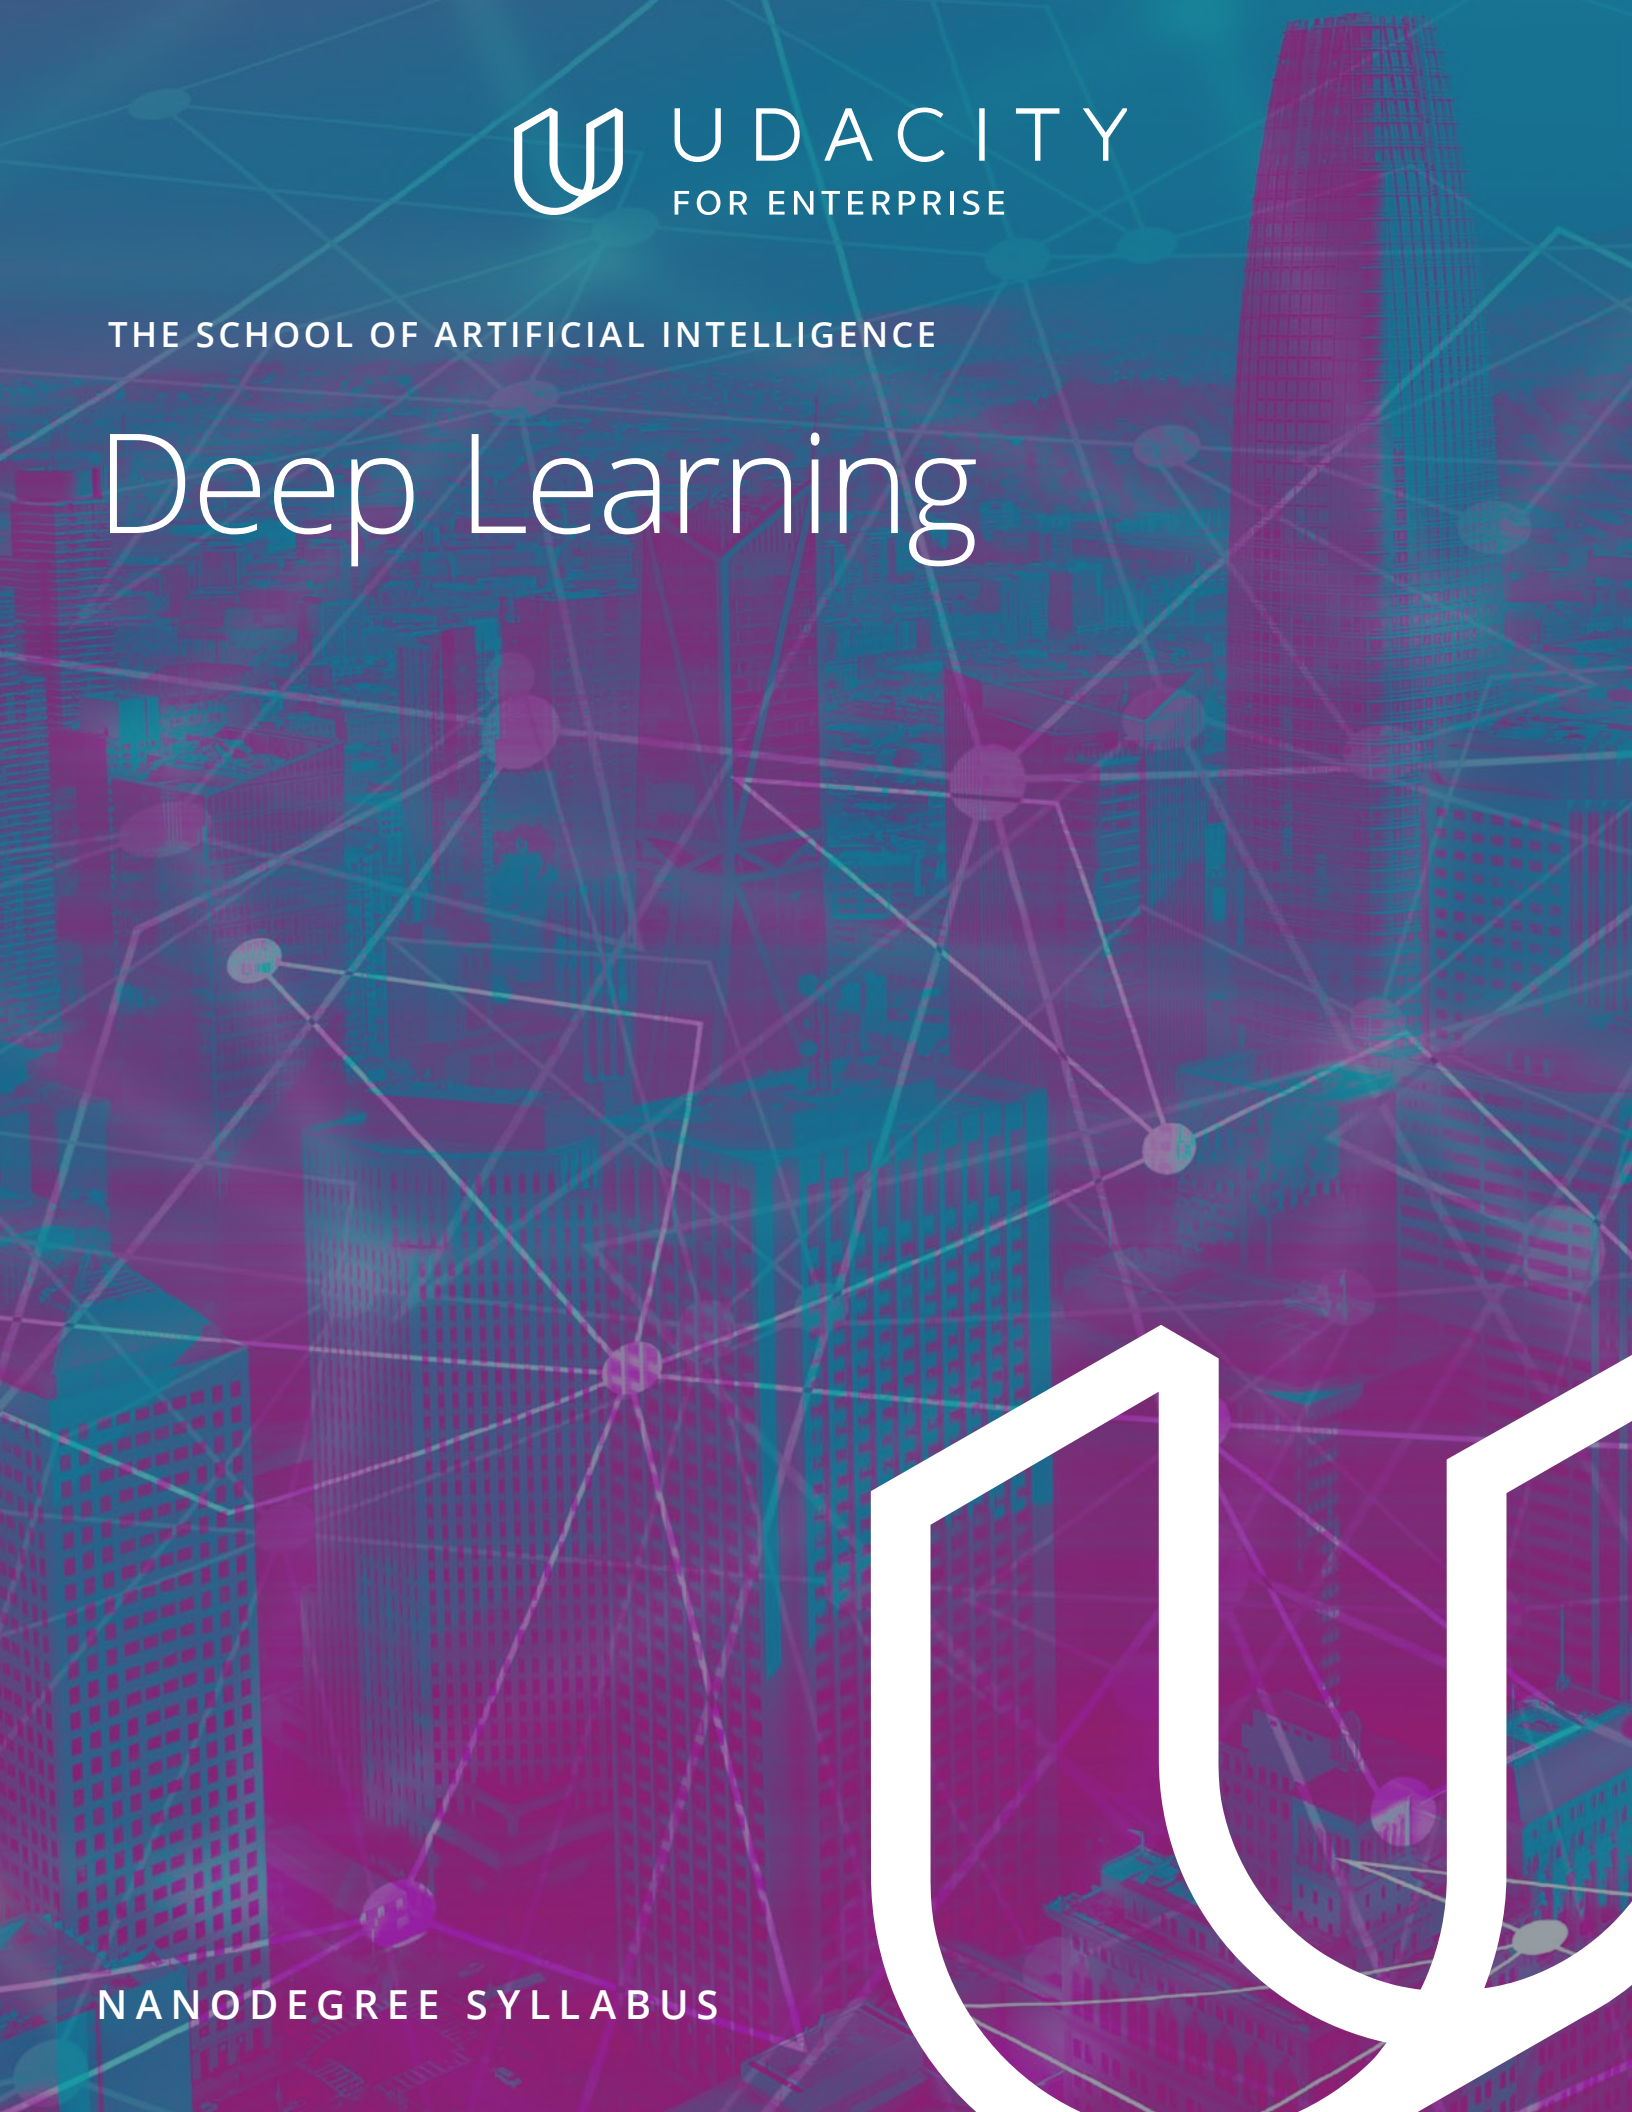
UDACITY
FOR ENTERPRISE
THE SCHOOL OF ARTIFICIAL INTELLIGENCE
Deep Learning
NANODEGREE SYLLABUS

Overview

This Nanodegree is Built in Partnership With

aws facebook Artificial Intelligence

The Deep Learning Nanodegree program offers a solid introduction to the world of artificial intelligence. This program is an ideal point-of-entry for anyone interested in this transformational technology.

In this program, learners master fundamentals that enable them to go further in the field and join the next generation of deep learning talent that will help define a beneficial, new, AI-powered future for our world. Learners study cutting-edge topics such as Neural Networks, Convolutional Neural Networks, Recurrent Neural Networks, Generative Adversarial Networks, and Network Deployment, and build projects in PyTorch and NumPy.

Learners receive instruction from authorities such as Ian Goodfellow and Jun-Yan Zhu, inventors of types of generative adversarial networks, as well as AI experts, Sebastian Thrun and Andrew Trask.

The program is comprised of 5 courses and 5 projects. Each project presents learners an opportunity to demonstrate mastery of newly acquired skills.

Program Information

TIME
4 months
Study 10 hours/week

LEVEL
Practitioner

PREREQUISITES
This program has been created specifically for learners familiar with machine learning, AI, and/or deep learning, and who have a working knowledge of Python.

HARDWARE/SOFTWARE REQUIRED
Computer running OS X or Windows

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

Our Classroom Experience

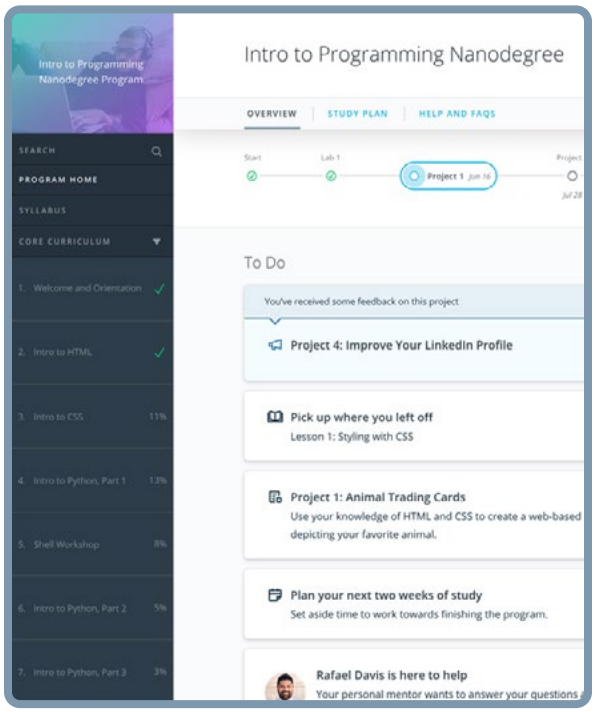

REAL-WORLD PROJECTS

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.

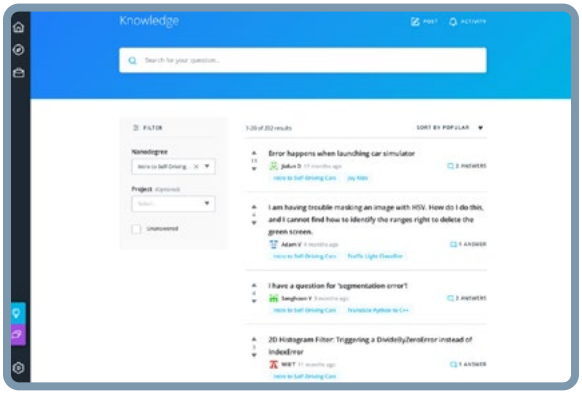

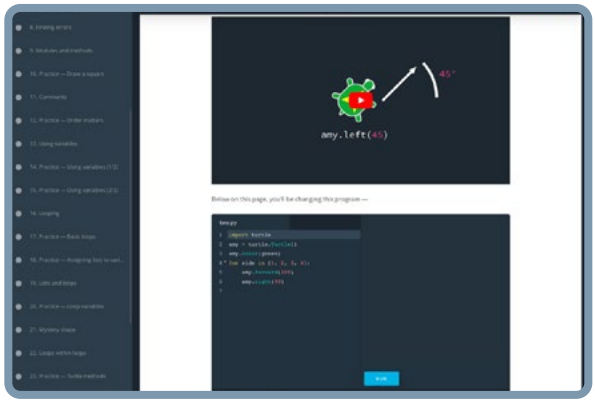

CUSTOM STUDY PLANS

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

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.

Luis Serrano

INSTRUCTOR

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.

Cezanne Camacho

CURRICULUM LEAD

Cezanne is a computer vision expert with a Master's in Electrical Engineering from Stanford University. As a former genomics and biomedical imaging researcher, she's applied computer vision and deep learning to medical diagnostics.

Alexis Cook

INSTRUCTOR

Alexis is an applied mathematician with a Masters in computer science from Brown University and a Masters in applied mathematics from the University of Michigan. She was formerly a National Science Foundation Graduate Research Fellow.

Jennifer Staab

INSTRUCTOR

Jennifer has 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.

Daniel Jiang

INSTRUCTOR

Daniel is a machine learning engineer who studied computer science at the University of California, Berkeley. He has worked on machine learning research at a variety of industry and academic groups.

Sean Carrell

INSTRUCTOR

Sean Carrell is a former research mathematician specializing in Algebraic Combinatorics. He completed his PhD and Postdoctoral Fellowship at the University of Waterloo, Canada.

Ortal Arel

INSTRUCTOR

Ortal Arel is a former computer engineering professor. She holds a PhD in Computer Engineering from the University of Tennessee. Her doctoral research work was in the area of applied cryptography.

Jay Alammar

INSTRUCTOR

Jay has a degree in computer science, loves visualizing machine learning concepts, and is the Investment Principal at STV, a $500 million venture capital fund focused on high-technology startups.

Course 1: Neural Networks

Learn neural networks basics, and build your first network with Python and Numpy. Use modern deep learning frameworks (NumPy, PyTorch) to build multi-layer neural networks, and analyze real data.

Project: Predicting Bike-Sharing Patterns

Learn neural networks basics, and build your first network with Python and NumPy. You'll define and train a multi-layer neural network, and use it to analyze real data. In this project, you will build and train neural networks from scratch to predict the number of bike-share users on a given day.

LESSON TITLE	LEARNING OUTCOME
INTRODUCTION TO NEURAL NETWORKS	In this lesson, you will learn solid foundations on deep learning and neural networks. You'll also implement gradient descent and backpropagation in Python.
IMPLEMENTING GRADIENT DESCENT	Mat and Luis will introduce you to a different error function and guide you through implementing gradient descent using NumPy matrix multiplication.
TRAINING NEURAL NETWORKS	Now that you are familiar with neural networks, this lesson introduces several techniques to improve network training. Learn how to prevent overfitting of training data and best practices for minimizing the error of a network.
SENTIMENT ANALYSIS	In this lesson, Andrew Trask, the author of Grokking Deep Learning, will show you how to define and train a neural networks for sentiment analysis (identifying and categorizing opinions expressed in text).
DEEP LEARNING WITH PYTORCH	Learn how to use PyTorch for building and testing deep learning models.

Course 2: Convolutional Neural Networks

Learn how to build convolutional networks and use them to classify images (faces, melanomas, etc.) based on objects that appear in them. Use these networks to learn data compression and image de-noising.

Project: Landmark Classification & Tagging for Social Media

In this project, you will apply the skills you have acquired in the course to build a landmark classifier. Photo-sharing services or photo-storage services may use landmark classification to automatically tag photos with relevant hashtags or location markers. This type of functionality could be especially important when photo location metadata is not available, which could happen when a photo is taken without metadata (e.g., phone was on airplane mode, camera was old and without GPS) or if a photo has had its metadata scrubbed. In the project, you will go through a machine learning design process end-to-end: performing data preprocessing and augmentation, designing your own CNN from scratch, and training and saving your best CNN model. You will also use transfer learning and compare your transfer-learned model with your from-scratch CNN.

LESSON TITLE	LEARNING OUTCOME
CLOUD COMPUTING	Take advantage of Amazon's GPUs to train your neural network faster. In this lesson, you'll setup an instance on AWS and train a neural network on a GPU.
CONVOLUTIONAL NEURAL NETWORK	Alexis and Cezanne explain how Convolutional Neural Networks can be used to identify patterns in images and how they help us dramatically improve performance in image classification tasks.
CNNs IN PYTORCH	In this lesson, you'll learn how to find good initial weights for a neural network. Having good initial weights often allows a neural network to arrive at an optimal solution, faster than without initialization.
AUTOENCODERS	Autoencoders are neural networks used for data compression, image denoising, and dimensionality reduction. Here, you'll build autoencoders using PyTorch.
TRANSFER LEARNING IN PYTORCH	Most people don't train their own networks on massive datasets. In this lesson, you'll learn how to finetune and use a pretrained network and apply it to a new task using transfer learning.
DEEP LEARNING FOR CANCER DETECTION	In this lesson, Sebastian Thrun teaches us about his groundbreaking work detecting skin cancer with Convolutional Neural Networks.

Course 3: Recurrent Neural Networks

Build your own recurrent networks and long short-term memory networks with NumPy and PyTorch; perform sentiment analysis and generate new text. Use recurrent networks to generate new text from TV scripts.

Project: Generate TV Scripts

In this project, you will build your own Recurrent Networks and Long Short-Term Memory Networks with PyTorch. You'll perform sentiment analysis and generate new text, and use recurrent networks to generate new text that resembles a training set of TV scripts.

LESSON TITLE	LEARNING OUTCOME
RECURRENT NEURAL NETWORKS	Ortal will introduce Recurrent Neural Networks (RNNs), which are machine learning models that are able to recognize and act on sequences of inputs.
LONG SHORT-TERM MEMORY NETWORK	Luis explains Long Short-Term Memory Networks (LSTM), and similar architectures that form a memory about a sequence of inputs, over time.
IMPLEMENTATION OF RNN & LSTM	Train recurrent neural networks to generate new characters, words, and bodies of text.
HYPERPARAMETERS	In this lesson, we'll look at a number of different hyperparameters that are important for our deep learning work, such as learning rates. We'll discuss starting values and intuitions for tuning each hyperparameter.
EMBEDDINGS & WORD2VEC	In this lesson, you'll learn about embeddings in neural networks by implementing a word2vec model that converts words into a representative vector of numerical values.
SENTIMENT PREDICTION RNN	In this lesson, you'll learn to implement a recurrent neural network for predicting sentiment. This is intended to give you more experience building RNNs.

Course 4: Generative Adversarial Networks

Learn to understand and implement the DCGAN model to simulate realistic images, with Ian Goodfellow, the inventor of GANS (generative adversarial networks), and Jun-Yan Zhu, the creator of CycleGANs.

Project: Generate Faces

Learn to understand Generative Adversarial Networks with the model's inventor, Ian Goodfellow. Then, apply what you've learned in this project and implement a Deep Convolutional GAN. This DCGAN is made of a pair of multi-layer neural networks that compete against each other until one learns to generate realistic images of faces.

LESSON TITLE	LEARNING OUTCOME
GENERATIVE ADVERSARIAL NETWORK	Ian Goodfellow, the inventor of GANs, introduces you to these exciting models. You'll also implement your own GAN on a simple dataset.
DEEP CONVOLUTIONAL GANs	Implement a Deep Convolutional GAN to generate complex, color images of house numbers.
PIX2PIX & CYCLEGAN	Jun-Yan Zhu and Cezanne lead you through a CycleGAN formulation that can learn from unlabeled sets of images.

Course 5: Updating a Model

In this project, you will train and deploy your own PyTorch sentiment analysis model using Amazon SageMaker on AWS. This model will be trained to do sentiment analysis on movie reviews (positive or negative reviews). You'll build the model, deploy it, and create a gateway for accessing this model from a website.

Project: Deploy a Sentiment Analysis Model

In this project, you will train and deploy your own PyTorch sentiment analysis model using Amazon SageMaker on AWS. This model will be trained to do sentiment analysis on movie reviews (positive or negative reviews). You'll build the model, deploy it, and create a gateway for accessing this model from a website.

LESSON TITLE	LEARNING OUTCOME
INTRODUCTION TO DEPLOYMENT	Learn where cloud deployment is used in industry and about various methods for deployment (websites, apps, etc.). Become familiar with cloud deployment terminology.
DEPLOY A MODEL	Deploy a model using Amazon SageMaker and learn to apply built-in algorithms, like XGBoost, to a variety of tasks.
CUSTOM MODELS & WEB HOSTING	In this lesson, you'll train and deploy your own PyTorch model. Then, see how to define a gateway using SageMaker to allow for outside-access to your model. See how your model responds to user input.
MODEL MONITORING	In this lesson, learn how to interpret log messages and monitor the behavior of your model over time. See how to implement an A/B test, in SageMaker, to evaluate the performance of two different models.
UPDATING A MODEL	Developing a machine learning model is an iterative process. Learn how to look at indicators like data distribution to see if you should update a model.

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.

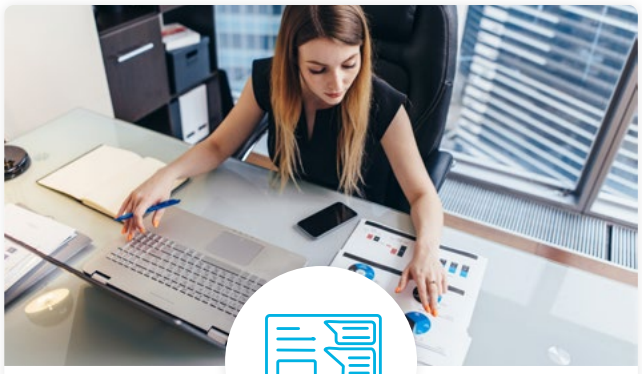

Dashboard & Progress Reports

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.

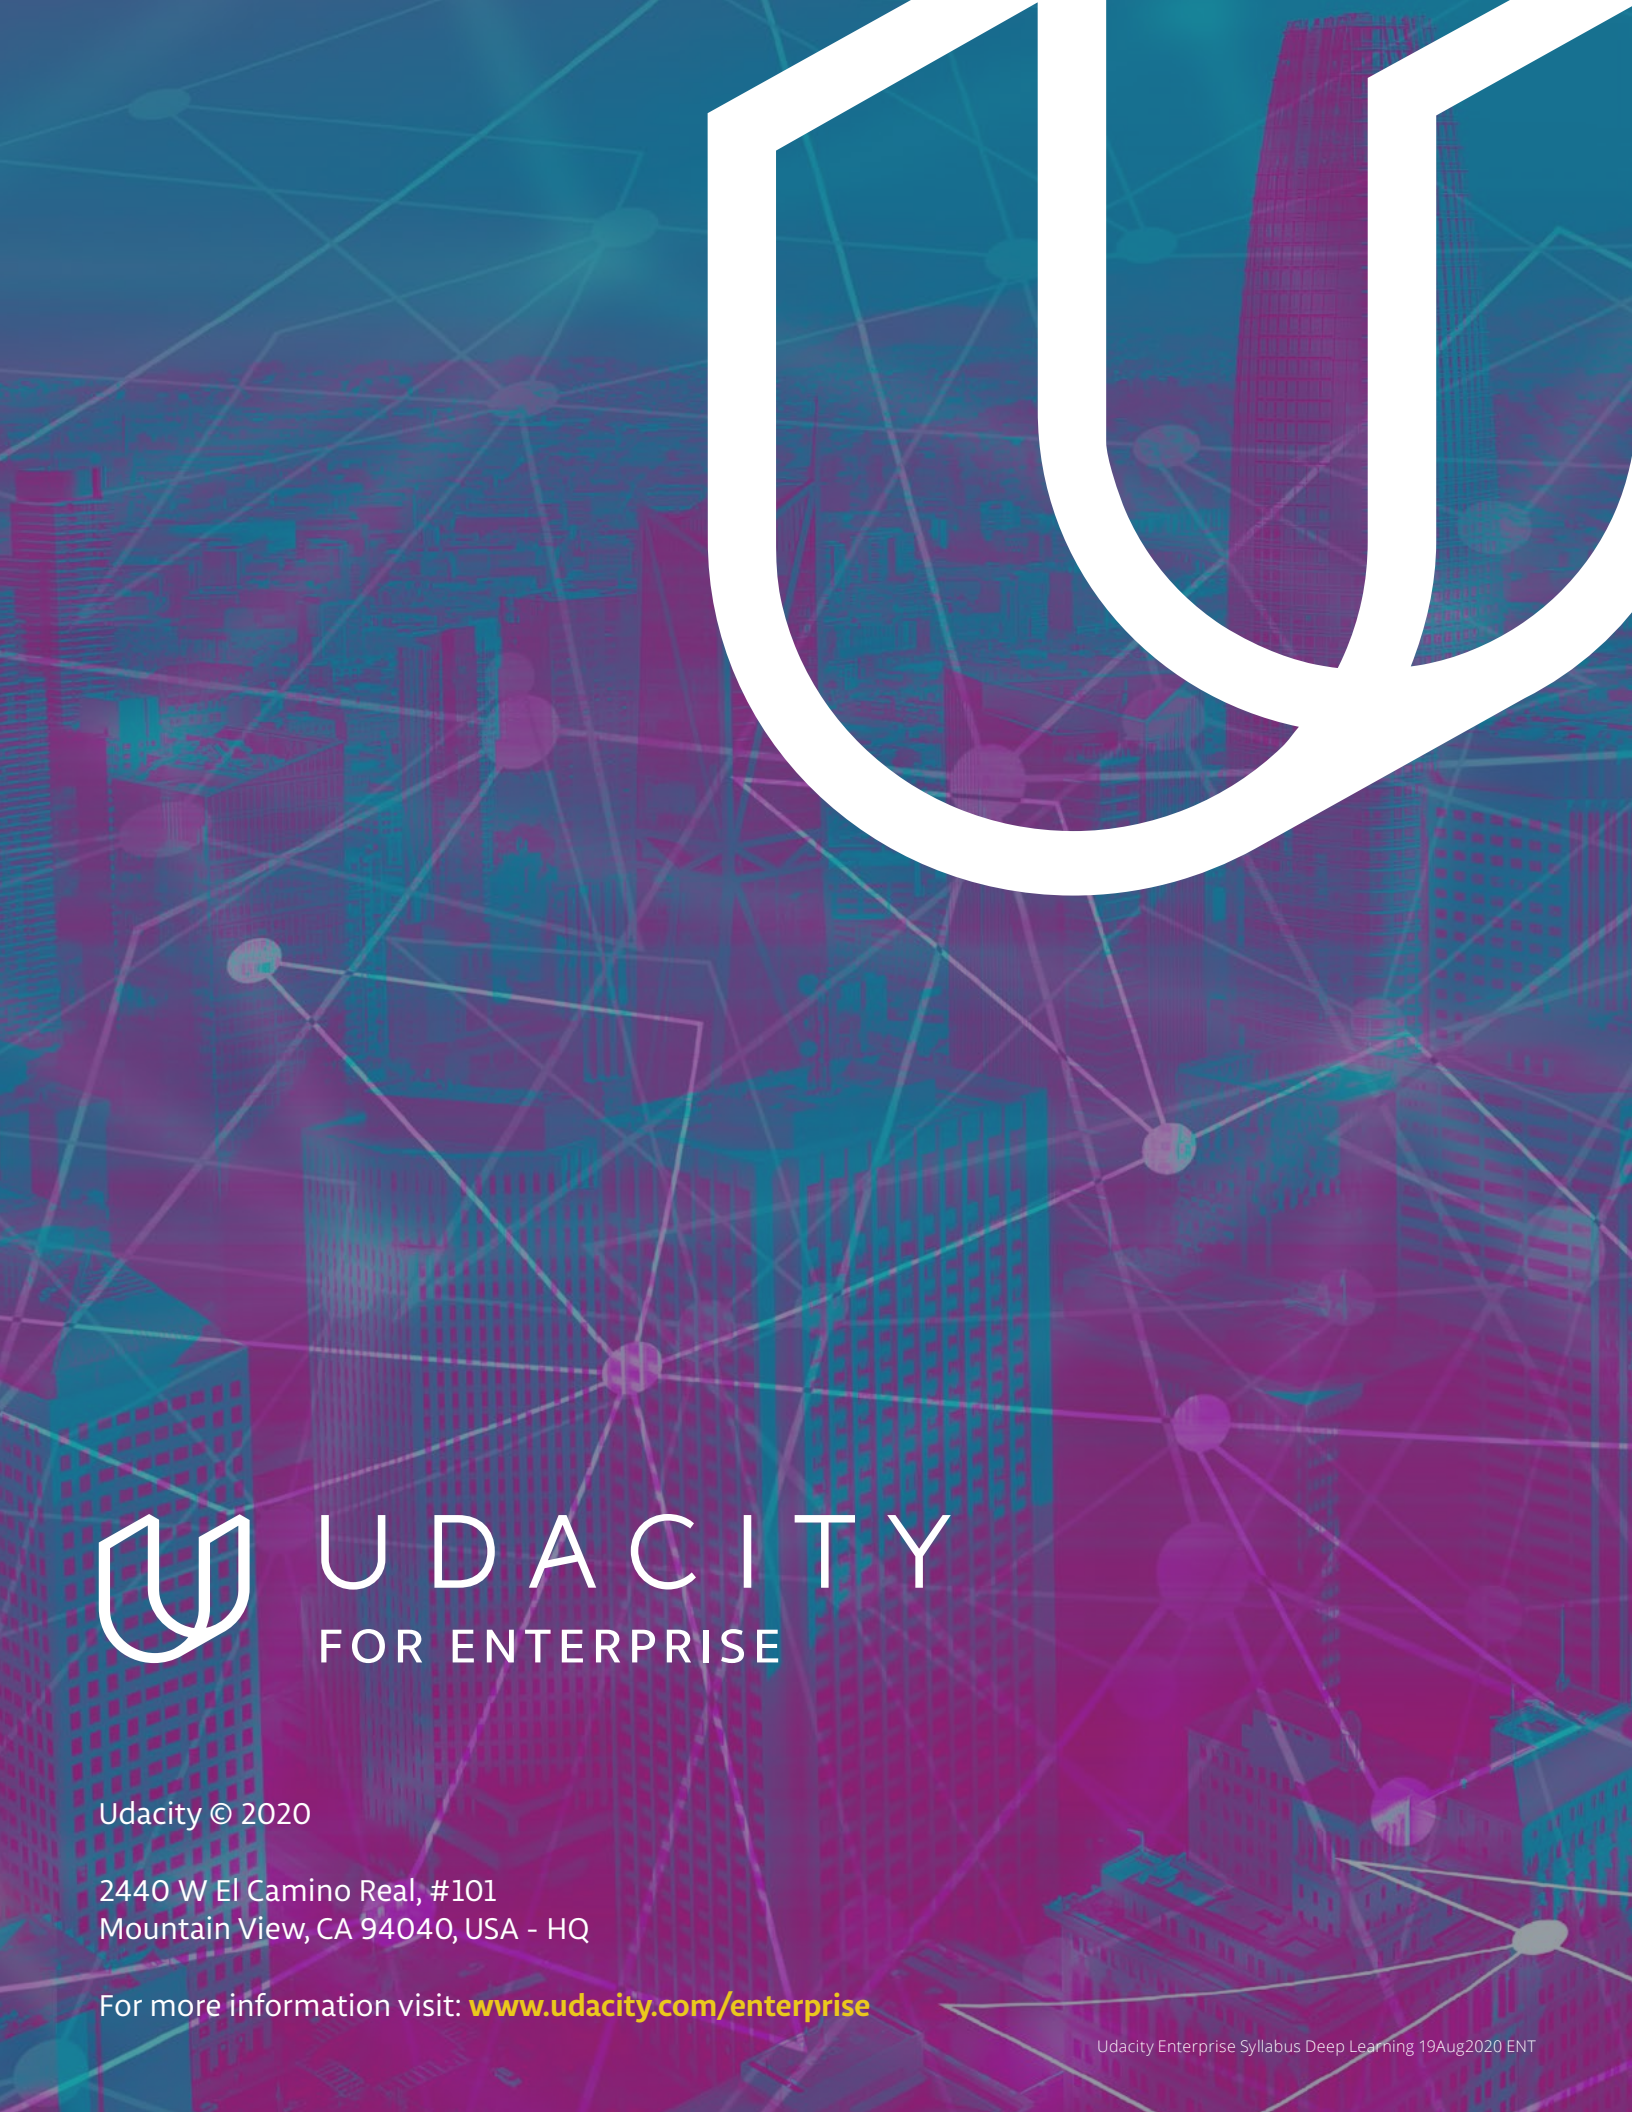

2440 W El Camino Real, #101
Mountain View, CA 94040, USA - HQ

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