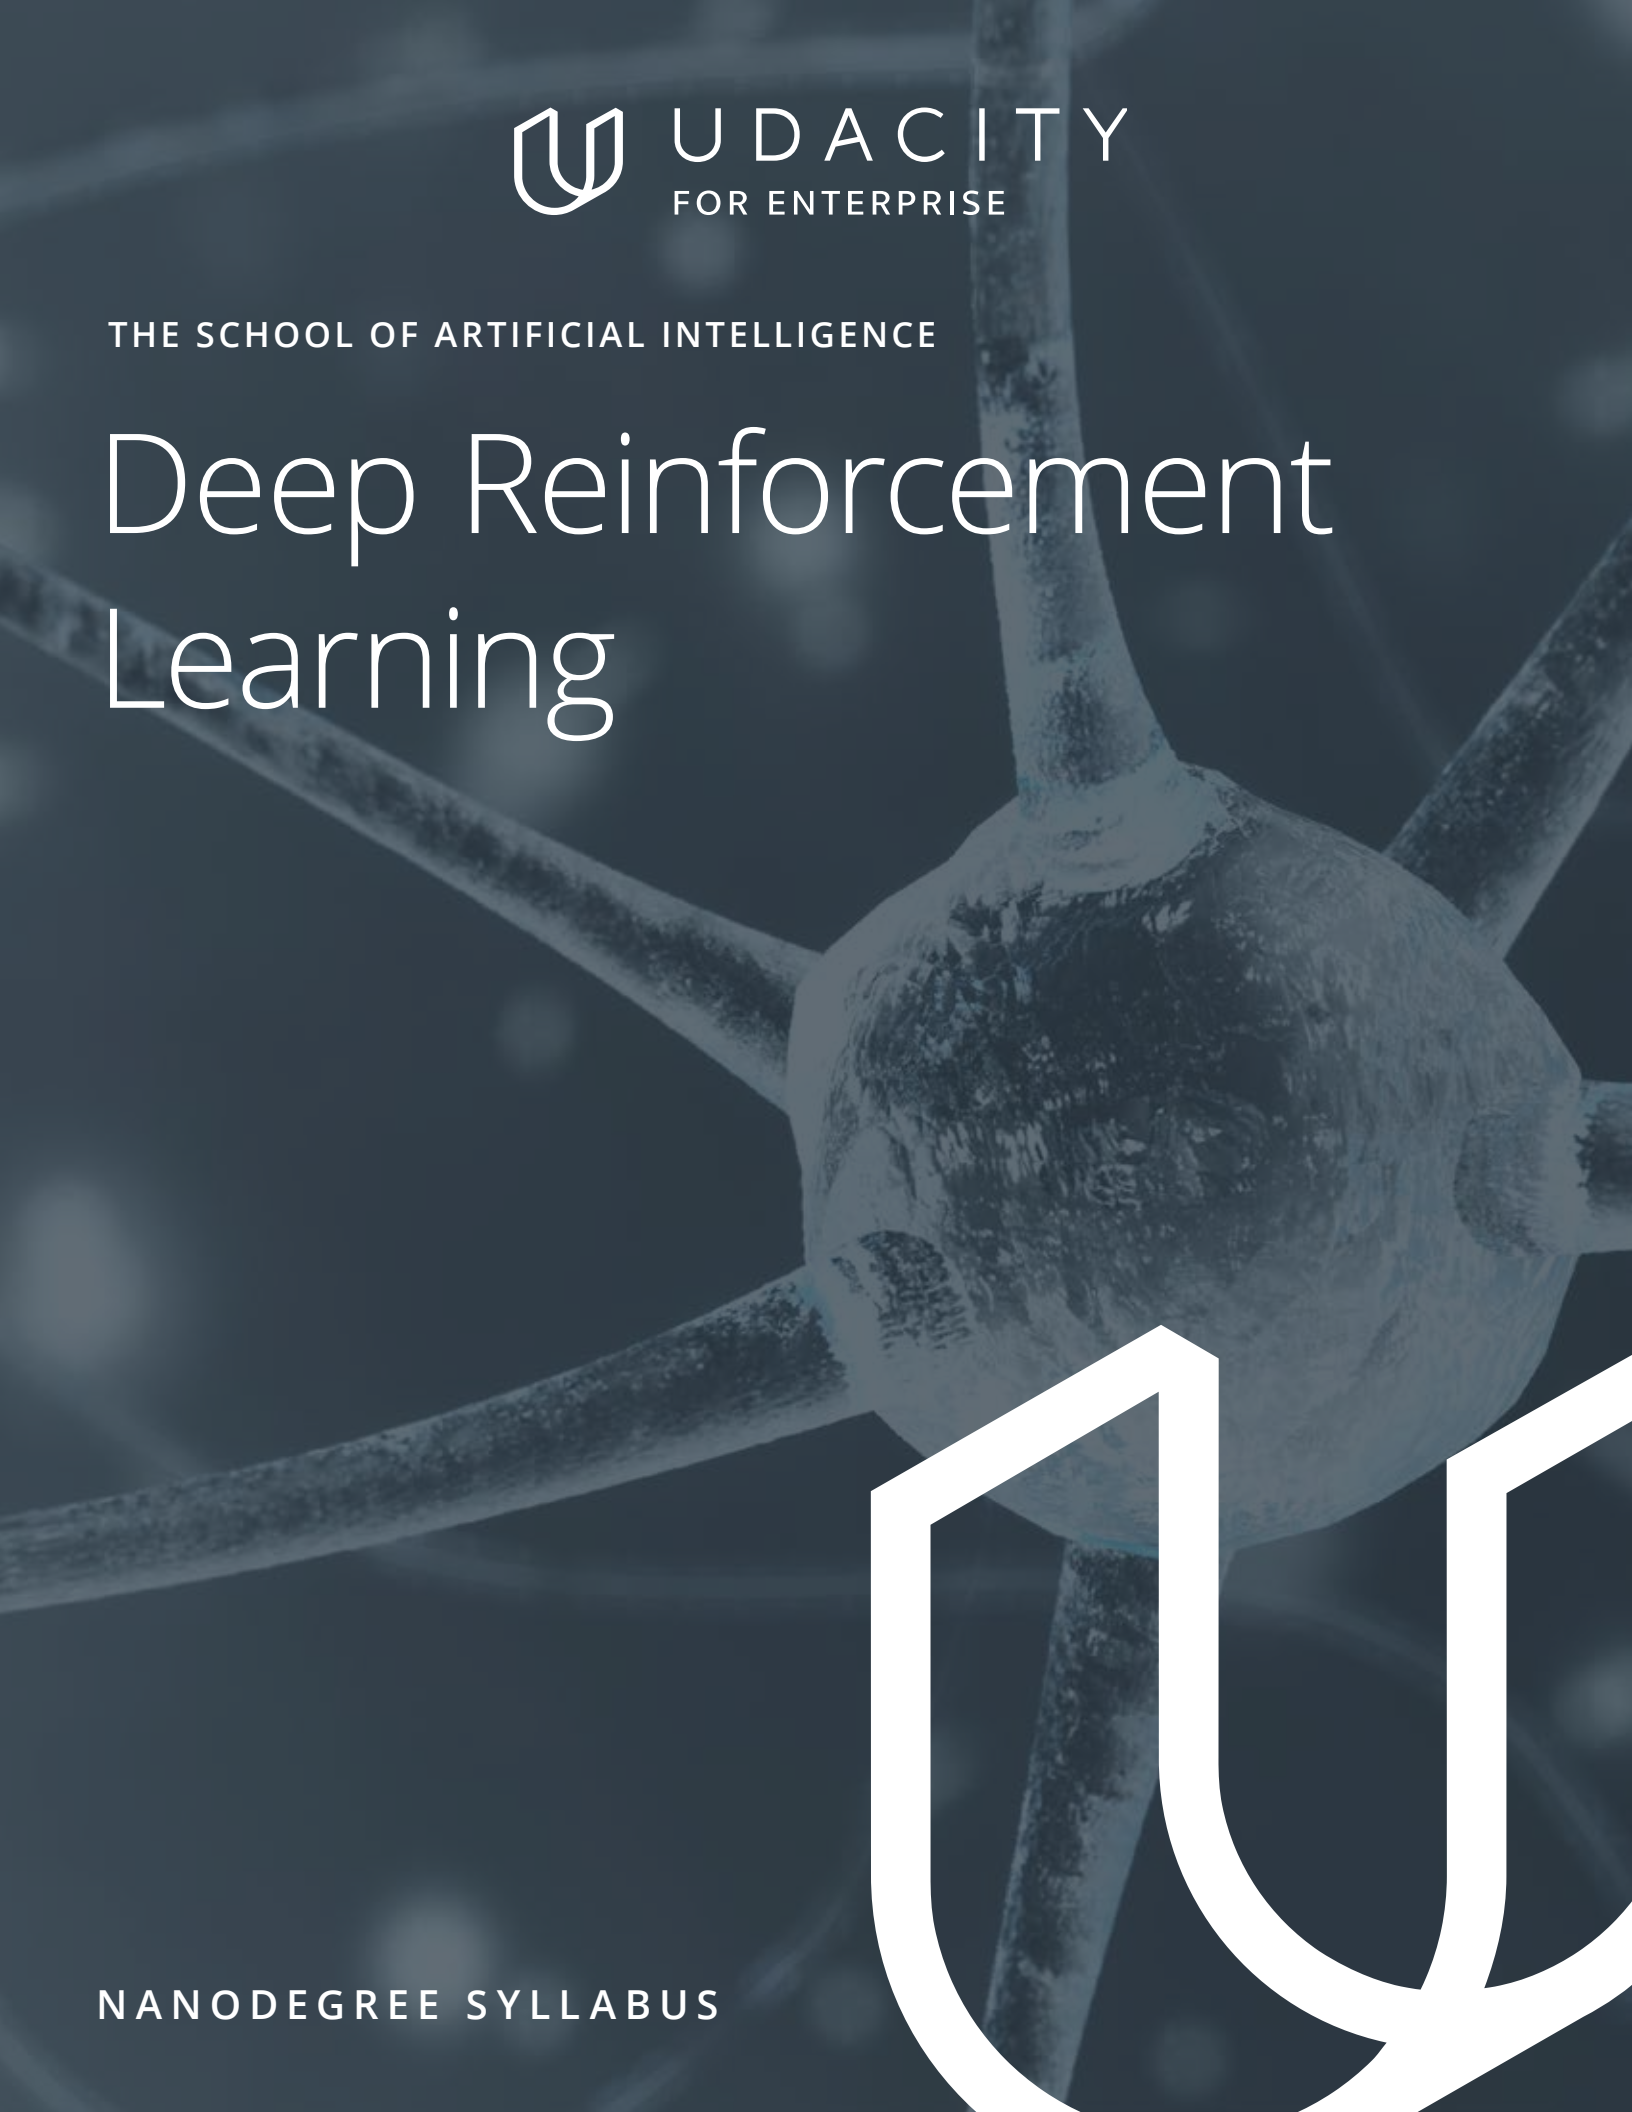

UDACITY
FOR ENTERPRISE

THE SCHOOL OF ARTIFICIAL INTELLIGENCE

Deep Reinforcement Learning

NANODEGREE SYLLABUS

Overview

This Nanodegree is Built in Partnership With

The Deep Reinforcement Learning Nanodegree program is designed to enhance your existing machine learning and deep learning skills with the addition of reinforcement learning theory and programming techniques.

This program will grow your deep learning and reinforcement learning expertise, and give you the skills you need to understand the most recent advancements in deep reinforcement learning, and build and implement your own algorithms.

The term is comprised of 4 courses and 3 projects, which are described in detail in this syllabus.

Building a project is one of the best ways to demonstrate the skills you've learned, and each project will contribute to a personal, impressive and professional portfolio that will showcase your mastery of reinforcement learning and deep learning techniques.

Program Information

TIME
4 months/study 10 hrs/week

LEVEL
Specialist

PREREQUISITES
- Intermediate to advanced Python experience.
- Familiarity with object-oriented programming.
- Read and understand code.
- Understanding of probability and statistics.
- Intermediate knowledge of machine learning techniques.
- Ability to describe backpropagation, and knowledge of neural network architectures (like a CNN for image classification).
- Experience with a deep learning framework like TensorFlow, Keras, or PyTorch.

HARDWARE/SOFTWARE REQUIRED
Computer running a 64-bit operating system with at least 8GB of RAM, along with administrator account permissions sufficient to install programs including Anaconda with Python 3.6 and supporting packages.

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

Our Classroom Experience

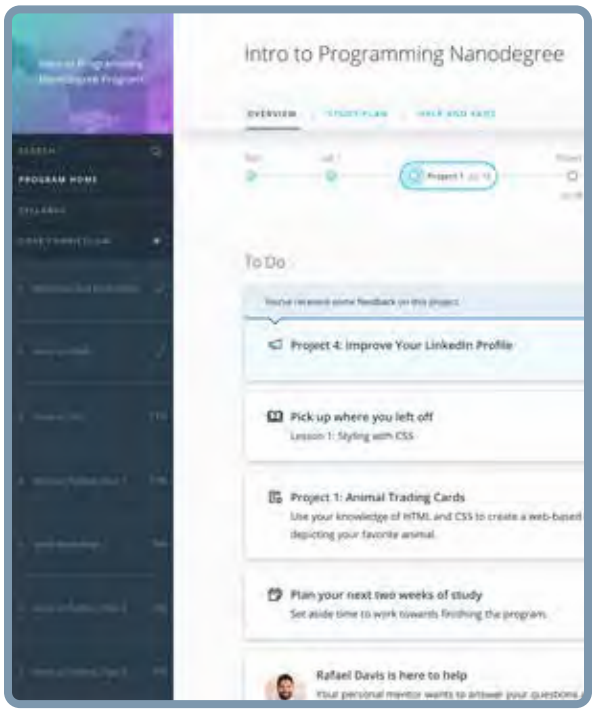

REAL-WORLD PROJECTS

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

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.

Arpan Chakraborty

INSTRUCTOR

Arpan is a computer scientist with a PhD from North Carolina State University. He teaches at Georgia Tech (within the Masters in Computer Science program), and is a coauthor of the book Practical Graph Mining with R.

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.

Learn with the Best

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.

Dana Sheahan

CONTENT DEVELOPER

Dana is an electrical engineer with a Masters in Computer Science from Georgia Tech. Her work experience includes software development for embedded systems in the Automotive Group at Motorola, where she was awarded a patent for an onboard operating system.

Chhavi Yadav

CONTENT DEVELOPER

Chhavi is a Computer Science graduate student at New York University, where she researches machine learning algorithms. She is also an electronics engineer and has worked on wireless systems.

Juan Delgado

CONTENT DEVELOPER

Juan is a computational physicist with a Masters in Astronomy. He is finishing his PhD in Biophysics. He previously worked at NASA developing space instruments and writing software to analyze large amounts of scientific data using machine learning techniques.

Miguel Morales

CONTENT DEVELOPER

Miguel is a software engineer at Lockheed Martin. He earned a Masters in Computer Science at Georgia Tech and is an Instructional Associate for the Reinforcement Learning and Decision Making course. He's the author of Grokking Deep Reinforcement Learning.

Nanodegree Program Overview

Course 1: Foundations of Reinforcement Learning

Master the fundamentals of reinforcement learning by writing your own implementations of many classical solution methods.

LESSON TITLE	LEARNING OUTCOME
INTRODUCTION TO RL	A friendly introduction to reinforcement learning.
The RL FRAMEWORK: THE PROBLEM	Learn how to define Markov Decision Processes to solve real-world problems.
THE RL FRAMEWORK: THE SOLUTION	Learn about policies and value functions. Derive the Bellman equations.
DYNAMIC PROGRAMMING	Write your own implementations of iterative policy evaluation, policy improvement, policy iteration, and value iteration.
MONTE CARLO METHODS	Implement classic Monte Carlo prediction and control methods. Learn about greedy and epsilon-greedy policies. Explore solutions to the Exploration-Exploitation Dilemma.
TEMPORAL-DIFFERENCE METHODS	Learn the difference between the Sarsa, Q-Learning, and Expected Sarsa algorithms.
SOLVE OPENAI GYM'S TAXI-V2 TASK	Design your own algorithm to solve a classical problem from the research community.
RL IN CONTINUOUS SPACES	Learn how to adapt traditional algorithms to work with continuous spaces.

Course 2: Value-Based Methods

Apply deep learning architectures to reinforcement learning tasks. Train your own agent that navigates a virtual world from sensory data.

Project: Navigation

Leverage neural networks to train an agent that learns intelligent behaviors from sensory data.

LESSON TITLE	LEARNING OUTCOME
DEEP LEARNING IN PYTORCH	Learn how to build and train neural networks and convolutional neural networks in PyTorch.
DEEP Q-LEARNING	Extend value-based reinforcement learning methods to complex problems using deep neural networks. Learn how to implement a Deep Q-Network (DQN), along with Double-DQN, Dueling-DQN, and Prioritized Replay.
DEEP RL FOR ROBOTICS	Learn from experts at NVIDIA how to use value-based methods in real-world robotics.

Project Example: Navigation

Video still from completed learner-submitted Deep Reinforcement Learning course project where learners train an agent to navigate (and collect bananas) in a large, square world.

Course 3: Policy-Based Methods

Learn the theory behind evolutionary algorithms and policy-gradient methods. Design your own algorithm to train a simulated robotic arm to reach target locations.

Project: Continuous Control

Train a robotic arm to reach target locations, or train a four-legged virtual creature to walk.

LESSON TITLE	LEARNING OUTCOME
INTRODUCTION TO POLICY-BASED METHODS	Learn the theory behind evolutionary algorithms, stochastic policy search, and the REINFORCE algorithm. Learn how to apply the algorithms to solve a classical control problem.
IMPROVING POLICY GRADIENT METHODS	Learn about techniques such as Generalized Advantage Estimation (GAE) for lowering the variance of policy gradient methods. Explore policy optimization methods such as Trust Region Policy Optimization (TRPO) and Proximal Policy Optimization (PPO).
ACTOR-CRITIC METHODS	Study cutting-edge algorithms such as Deep Deterministic Policy Gradients (DDPG).
DEEP RL FOR FINANCIAL TRADING	Learn from experts at NVIDIA how to use actor-critic methods to generate optimal financial trading strategies.

Project Example: Continuous Control

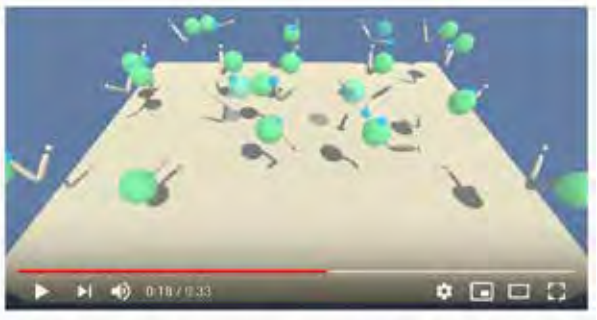

Udacity Deep Reinforcement Learning project 2 - After Training

Video still from completed learner-submitted Deep Reinforcement Learning course project where learners work with the Reacher environment to train a double-jointed arm to move to target locations, with the goal of the agent maintaining its position at the target location for as many time steps as possible.

Course 4: Multi-Agent Reinforcement Learning

Learn how to apply reinforcement learning methods to applications that involve multiple, interacting agents. These techniques are used in a variety of applications, such as the coordination of autonomous vehicles.

Project — Collaboration and Competition

Train a system of agents to demonstrate collaboration or cooperation on a complex task.

LESSON TITLE	LEARNING OUTCOME
INTRODUCTION TO MULTI-AGENT RL	Learn how to define Markov games to specify a reinforcement learning task with multiple agents. Explore how to train agents in collaborative and competitive settings.
CASE STUDY: ALPHAZERO	Master the skills behind DeepMind's AlphaZero.

Project Example — Collaboration and Competition

Video still from completed learner-submitted Deep Reinforcement Learning course project where learners work with the Tennis environment to train two agents within an observation space of 8 variables, to successfully keep the ball in play for a determined average score.

Our Nanodegree Programs Include:

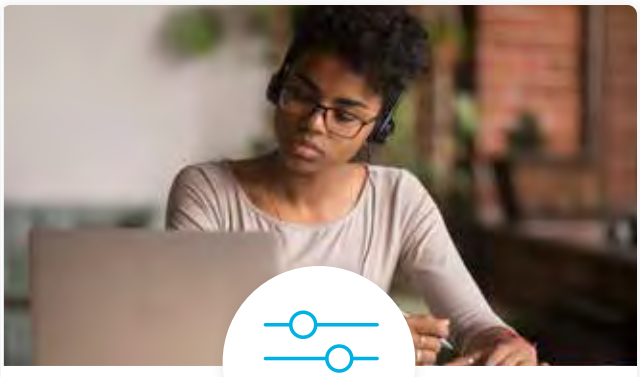

Pre-Assessments

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