



THE SCHOOL OF ARTIFICIAL INTELLIGENCE

Computer Vision



NANODEGREE SYLLABUS

Overview

This Nanodegree is Built in Partnership With



This program is designed to enhance your existing machine learning and deep learning skills with the addition of computer vision theory and programming techniques. These computer vision skills can be applied to various applications such as image and video processing, autonomous vehicle navigation, medical diagnostics, smartphone apps, and much more.

This program will not prepare you for a specific career or role, rather, it will grow your deep learning and computer vision expertise, and give you the skills you need to start applying computer vision techniques to real-world challenges and applications.

The term is comprised of 3 courses and 3 projects. Each project will contribute to an impressive professional portfolio that showcases your mastery of computer vision and deep learning techniques.

Program Information



TIME

3 months
Study 10 hours/week



LEVEL

Specialist



PREREQUISITES

Significant experience with Python, and entry-level experience with probability and statistics, and deep learning architectures. Ability to write a class in Python and add comments to your code for others to read. Familiarity with the term “neural networks” and differential math that drives backpropagation.



HARDWARE/SOFTWARE REQUIRED


A 64-bit operating system with at least 8GB of RAM, along with administrator account permissions sufficient to install programs including Anaconda with Python 3.5 and supporting packages. Your network should allow secure connections to remote hosts (like SSH).



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


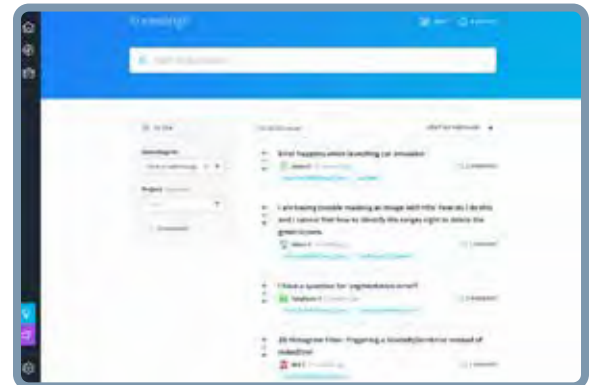
Learners leverage the power of community through a simple, yet powerful chat interface built within the classroom. Learner Hub connects learners with their technical mentor and fellow learners.

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



Sebastian Thrun

INSTRUCTOR

As the founder and president of Udacity, Sebastian's mission is to democratize education. He is also the founder of Google X, where he led projects including the Self-Driving Car, Google Glass, and more.



Cezanne Camacho

CURRICULUM LEAD

Cezanne is a machine learning educator with a Master's in Electrical Engineering from Stanford University. As a former researcher in genomics and biomedical imaging, she's applied machine learning to medical diagnostic applications.



Alexis Cook

CURRICULUM LEAD

Alexis is an applied mathematician with a Master's in Computer Science from Brown University and a Master's in Applied Mathematics from the University of Michigan. She was formerly a National Science Foundation Graduate Research Fellow.



Juan Delgado

CONTENT DEVELOPER

Juan is a computational physicist with a Master's 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.



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.



Ortal Arel

CURRICULUM LEAD

Ortal holds a PhD in Computer Engineering, and has been a professor and researcher in the field of applied cryptography. She has worked on design and analysis of intelligent algorithms for high-speed custom digital architectures.



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.



Course 1: Introduction to Computer Vision

Master computer vision and image processing essentials. Learn to extract important features from image data, and apply deep learning techniques to classification tasks.

Project

Facial Keypoint Detection

Use image processing techniques and deep learning techniques to detect faces in an image and find facial keypoints, such as the position of the eyes, nose, and mouth on a face. This project tests your knowledge of image processing and feature extraction techniques that allow you to programmatically represent different facial features. You'll also use your knowledge of deep learning techniques to program a convolutional neural network to recognize facial keypoints. Facial keypoints including points around the eyes, nose, and mouth on any face and are used in many applications, from facial tracking to emotion recognition.

LESSON TITLE

LEARNING OUTCOME

INTRODUCTION TO COMPUTER VISION

- Learn where computer vision techniques are used in industry.
- Prepare for the course ahead with a detailed topic overview.
- Start programming your own applications.

IMAGE REPRESENTATION AND ANALYSIS

- See how images are represented numerically.
- Implement color and geometric transforms image processing techniques.
- Program your own convolutional kernel for object edge-detection.

CONVOLUTIONAL NN LAYERS

- Learn about the layers of a deep convolutional neural network: convolutional, maxpooling, and fully-connected layers.
- Build an CNN-based image classifier in PyTorch.
- Learn about layer activation and feature visualization techniques.

FEATURES AND OBJECT RECOGNITION

- Learn why distinguishing features are important in pattern and object recognition tasks.
- Write code to extract information about an object's color and version 1.0 shape.
- Use features to identify areas on a face and to recognize the shape of a car or pedestrian on a road.

IMAGE SEGMENTATION

- Implement k-means clustering to break an image up into parts.
- Find the contours and edges of multiple objects in an image.
- Learn about background subtraction for video.

Nanodegree Program Overview

Course 2: Advanced Computer Vision and Deep Learning

Learn to apply deep learning architectures to computer vision tasks. Discover how to combine CNN and RNN networks to build an automatic image captioning application.

Project

Automatic Image Captioning

Combine CNN and RNN knowledge to build a deep learning model that produces captions given an input image. Image captioning requires that you create a complex deep learning model with two components: a CNN that transforms an input image into a set of features, and an RNN that turns those features into rich, descriptive language. In this project, you will implement these cutting-edge deep learning architectures.

LESSON TITLE

LEARNING OUTCOME

ADVANCED CNN ARCHITECTURES

- Learn about advances in CNN architectures.
- See how region-based CNN's, like Faster R-CNN, have allowed for fast, localized object recognition in images.
- Work with a YOLO/single shot object detection system.

RECURRENT NEURAL NETWORKS

- Learn how recurrent neural networks learn from ordered sequences of data.
- Implement an RNN for sequential text generation.
- Explore how memory can be incorporated into a deep learning model.
- Understand where RNN's are used in deep learning applications.

ATTENTION MECHANISMS

- Learn how attention allows models to focus on a specific piece of input data.
- Understand where attention is useful in natural language and computer vision applications.

IMAGE CAPTIONING

- Learn how to combine CNNs and RNNs to build a complex captioning model.
- Implement an LSTM for caption generation.
- Train a model to predict captions and understand a visual scene.



Project Example

Automatic Image Captioning

Image Captioning Project

In this project, I design and train a CNN-RNN (Convolutional Neural Network - Recurrent Neural Network) model for automatically generating image captions. The network is trained on the Microsoft Common Objects in COntext (MS COCO) dataset. The image captioning model is displayed below.

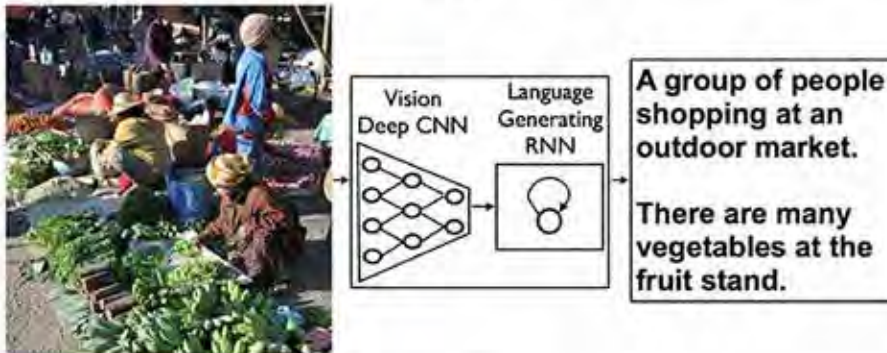


Image source

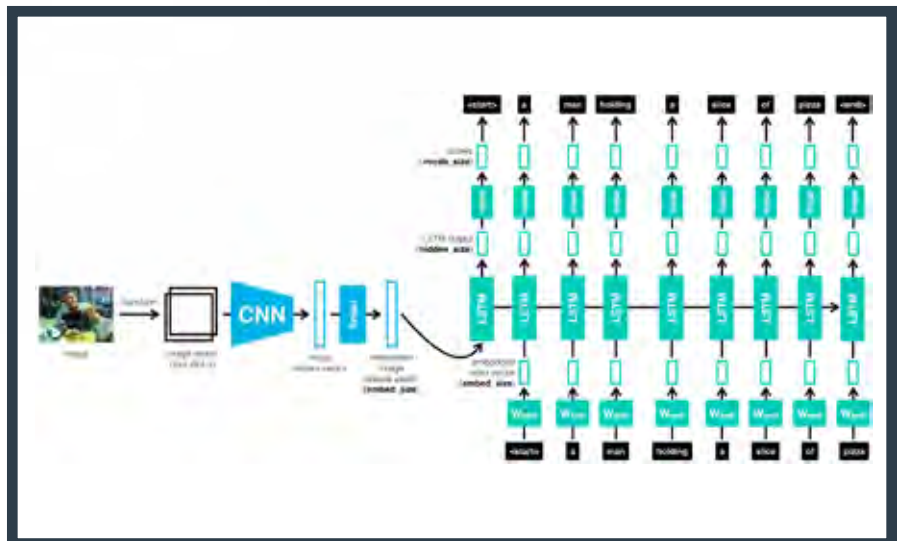
One good and one not so good sample made by my model:



A large propeller plane parked on the runway.



A cat is sitting on a toilet seat.



Actual Computer Vision Nanodegree course project material

Nanodegree Program Overview

Course 3: Object Tracking and Localization

Learn how to locate an object and track it over time. These techniques are used in a variety of moving systems, such as self-driving car navigation and drone flight.

Project

Landmark Detection and Tracking

Use feature detection and keypoint descriptors to build a map of the environment with SLAM (simultaneous localization and mapping). Implement a robust method for tracking an object over time, using elements of probability, motion models, and linear algebra. This project tests your knowledge of localization techniques that are widely used in autonomous vehicle navigation.

LESSON TITLE	LEARNING OUTCOME
OBJECT MOTION AND TRACKING	<ul style="list-style-type: none">• Learn how to programmatically track a single point over time.• Understand motion models that define object movement over time.• Learn how to analyze videos as sequences of individual image frames.
OPTICAL FLOW AND FEATURE MATCHING	<ul style="list-style-type: none">• Implement a method for tracking a set of unique features over time.• Learn how to match features from one image frame to another.• Track a moving car using optical flow.
ROBOT LOCALIZATION	<ul style="list-style-type: none">• Use Bayesian statistics to locate a robot in space.• Learn how sensor measurements can be used to safely navigate an environment.• Understand Gaussian uncertainty.• Implement a histogram filter for robot localization in Python.
GRAPH SLAM	<ul style="list-style-type: none">• Identify landmarks and build up a map of an environment.• Learn how to simultaneously localize an autonomous vehicle and create a map of landmarks.• Implement move and sense functions for a robotic vehicle.




Project Example

Landmark Detection and Tracking

Landmark Detection Project

In this project, I implement SLAM (Simultaneous Localization and Mapping) for a 2-dimensional world. Sensor and motion data gathered by a simulated robot is used to create a map of an environment. SLAM gives us a way to track the location of a robot in the world in real-time and identify the locations of landmarks such as buildings, trees, rocks, etc.

Last pose: (9.816799097356466, 42.351095108474404)



$$\mu = \Omega^{-1} \xi$$

Solutions


$$x_0 = -3$$

$$x_1 = 2$$

$$x_2 = 5$$

Actual Computer Vision Nanodegree course project material

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.



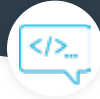
Vaibhav
UDACITY LEARNER

"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 © 2020

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

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