



THE SCHOOL OF ARTIFICIAL INTELLIGENCE

Intro to Machine Learning with Tensor Flow

NANODEGREE SYLLABUS



Overview

This Nanodegree is Built in Partnership With



The ultimate goal of the Intro to Machine Learning with TensorFlow Nanodegree program is to help learners acquire mastery of machine learning techniques such as data transformation and algorithms that can find patterns in data and apply machine learning algorithms to tasks of their own design.

A graduate of this program will be able to:

- Use Python and SQL to access and analyze data from several different data sources.
- Build predictive models using a variety of unsupervised and supervised machine learning techniques.
- Perform feature engineering to improve the performance of machine learning models.
- Optimize, tune, and improve algorithms according to specific metrics like accuracy and speed.
- Compare the performances of learned models using suitable metrics.

This program is comprised of 3 courses and 3 projects. Each project you build will be an opportunity to demonstrate what you've learned in the lessons. Your completed projects will become part of a portfolio that will showcase your mastery of machine learning techniques such as customer segmentation and image classification.

Program Information



TIME

3 months
Study 10 hours/week



LEVEL

Practitioner



PREREQUISITES

Intermediate Python programming knowledge, and basic understanding of probabilities and statistics.



HARDWARE/SOFTWARE REQUIRED


A 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.x 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



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.



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.



Dan Romuald Mbanga

INSTRUCTOR

Dan leads Amazon AI's Business Development efforts for Machine Learning Services. Day to day, he works with customers—from startups to enterprises—to ensure they are successful at building and deploying models on Amazon SageMaker.



Jennifer Staab

INSTRUCTOR

Jennifer has a PhD in Computer Science and a Masters 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.



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.



Josh Bernard

DATA SCIENTIST AT NERD WALLET

Josh has been sharing his passion for data for nearly a decade at all levels at university, and as Lead Science Instructor at Galvanize. He's utilized data science for work ranging from cancer research to process automation.



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.



Andrew Paster

INSTRUCTOR

Andrew has an engineering degree from Yale, and has used his data science skills to build a jewelry business from the ground up. He has additionally created courses for Udacity's Self-Driving Car Engineer Nanodegree program.



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.



Michael Virgo

INSTRUCTOR

Michael is a Content Developer at Udacity, and is currently pursuing a Masters in CS. After beginning his career in business, he utilized Udacity Nanodegree programs to build his technical skills, eventually becoming a Self-Driving Car Engineer at Udacity before switching roles to work on content.

Nanodegree Program Overview

Course 1: Supervised Learning

Learn about supervised learning, a common class of methods for model construction.

Project

Find Donors for CharityML

CharityML is a fictitious charity organization located in the heart of Silicon Valley that was established to provide financial support for people eager to learn machine learning. To expand their potential donor base, CharityML has decided to send letters to residents of California, but to only those most likely to donate to the charity. Your goal will be to evaluate and optimize several different supervised learning algorithms to determine which algorithm will provide the highest donation yield while under some marketing constraints.

LESSON TITLE

LEARNING OUTCOME

REGRESSION

- Learn the difference between Regression and Classification.
- Train a Linear Regression model to predict values.
- Learn to predict states using Logistic Regression.

PERCEPTRON ALGORITHMS

- Learn the definition of a perceptron as a building block for neural networks, and the perceptron algorithm for classification.

DECISION TREES

- Train Decision Trees to predict states.
- Use Entropy to build decision trees, recursively.

NAIVE BAYES

- Learn Bayes' rule, and apply it to predict cases of spam messages using the Naive Bayes algorithm.
- Train models using Bayesian Learning.
- Complete an exercise that uses Bayesian Learning for natural language processing.

SUPPORT VECTOR MACHINES

- Learn to train a Support Vector Machines to separate data, linearly.
- Use Kernel Methods in order to train SVMs on data that is not linearly separable.



Course 1: Supervised Learning, cont.

LESSON TITLE	LEARNING OUTCOME
ENSEMBLE OF LEARNERS	<ul style="list-style-type: none">• Build data visualizations for quantitative and categorical data.• Create pie, bar, line, scatter, histogram, and boxplot charts.• Build professional presentations.
EVALUATION METRICS	<ul style="list-style-type: none">• Learn about different metrics to measure model success.• Calculate accuracy, precision, and recall to measure the performance of your models.
TRAINING AND TUNING MODELS	<ul style="list-style-type: none">• Train and test models with Scikit-learn.• Choose the best model using evaluation techniques like cross-validation and grid search.



Nanodegree Program Overview

Course 2: Neural Networks

Learn the foundations of neural network design and training in TensorFlow.

Project

Create an Image Classifier

Implementing an image classification application using a deep neural network. This application will train a deep learning model on a dataset of images. It will then use the trained model to classify new images. You will develop your code in a Jupyter notebook to ensure your implementation works well. Key skills demonstrated include TensorFlow and neural networks, and model validation and evaluation.

LESSON TITLE

LEARNING OUTCOME

INTRODUCTION TO NEURAL NETWORKS

- Learn the foundations of deep learning and neural networks.
- Implement gradient descent and backpropagation in Python.

IMPLEMENTING GRADIENT DESCENT

- Implement gradient descent using NumPy matrix multiplication.

TRAINING NEURAL NETWORKS

- Learn several techniques to effectively train a neural network.
- Prevent overfitting of training data and learn best practices for minimizing the error of a network.

DEEP LEARNING WITH TENSORFLOW

- Learn how to use TensorFlow for building deep learning models.





Course 3: Unsupervised Learning

Learn to implement unsupervised learning methods for different kinds of problem domains.

Project

Create Customer Segments

Apply unsupervised learning techniques on product spending data collected for customers of a wholesale distributor in Lisbon, Portugal to identify customer segments hidden in the data. You will first explore and pre-process the data by scaling each product category and then identifying (and removing) unwanted outliers. With the cleaned customer spending data, you will apply PCA transformations to the data and implement clustering algorithms to segment the transformed customer data. Finally, you will compare the segmentation found with an additional labeling and consider ways this information could assist the wholesale distributor with future service changes. Key skills demonstrated include data cleaning, dimensionality reduction with PCA, and unsupervised clustering.

LESSON TITLE

LEARNING OUTCOME

CLUSTERING

- Learn the basics of clustering data.
- Cluster data with the K-means algorithm.

HIERARCHICAL AND DENSITY-BASED CLUSTERING

- Cluster data with Single Linkage Clustering.
- Cluster data with DBSCAN, a clustering method that captures the insight that clusters are dense group of points.



GAUSSIAN MIXTURE MODELS

- Cluster data with Gaussian Mixture Models.
- Optimize Gaussian Mixture Models with and Expectation Maximization.

DIMENSIONALITY REDUCTION

- Reduce the dimensionality of the data using principal component analysis and Independent Component Analysis.

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