

Artificial Intelligence Nanodegree and Specializations Syllabus



Contact Info

While going through the program, if you have questions about anything, you can reach us at support@udacity.com. For help from Udacity Mentors and your peers visit the Udacity Classroom.

Nanodegree Program Info

Version: 1.0.0

Length of Program: 260 Days*

** This is a self-paced program and the length is an estimation of total hours the average student may take to complete all required coursework, including lecture and project time. Actual hours may vary.*

Part 1: Introduction to Artificial Intelligence

In this Term, you'll learn the foundations of AI with Sebastian Thrun, Peter Norvig, and Thad Starner. We'll cover Game-Playing, Search, Optimization, Probabilistic AIs, and Hidden Markov Models.

Project: Solve a Sudoku with AI

Supporting Lessons

Lesson

Summary

Welcome to Artificial Intelligence

Welcome to **Term 1** of the Artificial Intelligence Nanodegree program!

Introduction to AI

An introduction to basic AI concepts with real world examples.

Applying AI to Sudoku

In this lesson, you'll get taste of the power of Artificial Intelligence by developing an algorithm to solve every Sudoku puzzle. Enjoy the fun of building your first AI agent and get coding!

Setting up with Anaconda

Get your environment set up using Anaconda, an extremely popular way to manage your environments and packages in python.

Project: Build a Game-Playing Agent

Supporting Lessons

Lesson

Summary

Introduction to Game Playing

In this lesson, you'll learn about how to build a Game Playing AI Agent. You'll focus on an agent that wins in the board game Isolation! You'll learn some of the seminal techniques in AI Game Playing including Adversarial Search and Minimax Trees.

Advanced Game Playing

In this lesson, you'll build a Game-Playing agent that defeats opponents in Isolation. Along the way, you'll learn about advanced Game-Playing techniques such as Iterative Deepening, Alpha-Beta Pruning, and Expectimax.

Project: Implement a Planning Search

Supporting Lessons

Lesson

Summary

Logic and Reasoning

In this lesson you'll learn to build intelligent systems that can reason using logic! This is in many ways one of the foundational pieces of intelligence - the ability to arrive at new conclusions from a given set of facts.

Planning

Explore how we can use logic and search to plan out complex itineraries. Many of these planning approaches are the same ones used to power Self-Driving Cars!

Project: Build a Sign Language Recognizer

Supporting Lessons

Lesson	Summary
Probability	Learn to model uncertainty in the real world using probability theory.
Bayes Nets	Learn to encode probability distributions using compact graphical models that enable efficient analysis.
Inference in Bayes Nets	In this lesson, you will learn about probabilistic inference using Bayes Nets, i.e. how to answer questions that you are interested in, given certain inputs.
Hidden Markov Models	Learn to process sequences and time-series data using Hidden Markov Models.

Part 2: Deep Learning and Applications

In this term, you'll learn the cutting edge advancements of AI and Deep Learning. You'll get the chance to apply Deep Learning on a variety of different topics including Computer Vision, Speech, and Natural Language Processing. We'll cover Convolutional Neural Networks, Recurrent Neural Networks, and other advanced models.

Project: Dog Breed Classifier

In this project, you will learn how to build a pipeline to process real-world, user-supplied images. Given an image of a dog, your algorithm will identify an estimate of the canine's breed.

Supporting Lessons

Lesson	Summary
Convolutional Neural Networks	Alexis explains the theory behind Convolutional Neural Networks and how they help us dramatically improve performance in image classification.

Project: Time Series Prediction and Text Generation

In this project you'll build RNNs that can generate sequences based on input data.

Supporting Lessons

Lesson	Summary
Recurrent Neural Networks	Jeremy explains Recurrent Neural Networks, and their cutting edge applications to text-based sequence generation
Long Short-Term Memory Networks (LSTM)	Luis explains Long Short-Term Memory Networks (LSTM), and similar architectures which have the benefits of preserving long term memory.
Implementing RNNs and LSTMs	In this lesson, Mat will review the concepts of RNNs and LSTMs, and then you'll see how a character-wise recurrent network is implemented in TensorFlow.
Hyperparameters	In this section, Jay will teach you about some important hyperparameters used for our deep learning work, including those used for Recurrent Neural Networks.
Sentiment Prediction with RNN	In this lesson you'll implement a sentiment prediction RNN

Part 3: Computer Vision

In this module, you will learn how to build intelligent systems that can see and understand the world using Computer Vision. You'll learn fundamental techniques for tasks like Object Recognition, Face Detection, Video Analysis, etc., and integrate classic methods with more modern Convolutional Neural Networks.

Project: CV Capstone Project

You'll apply what you've learned about the computer vision pipeline and build an end-to-end **facial** keypoint recognition system!

Supporting Lessons

Lesson	Summary
Image Representation and Analysis	In this section, you'll learn the fundamentals of computer vision: from how an image is formed to how to process, filter, and transform images!
Image Segmentation	In this section, you'll learn how to break an image up into segments and areas of interest using a variety of different algorithms!
Features and Object Recognition	Here, we discuss the end goal of many computer vision applications: feature extraction and object recognition. We'll cover why these are important applications and how to code them!

Part 4: Natural Language Processing

In this module, you will build end-to-end Natural Language Processing pipelines, starting from text processing, to feature extraction and modeling for different tasks such as Sentiment Analysis, Spam Detection and Machine Translation. You'll also learn how to design Recurrent Neural Networks for challenging NLP applications.

Project: Machine Translation

Apply the skills you've learnt in Natural Language Processing to the challenging and extremely rewarding task of Machine Translation. *_Bonne chance!_*

Supporting Lessons

Lesson	Summary
Natural Language Processing	An overview of how to build an end-to-end Natural Language Processing pipeline.
Text Processing	Learn to prepare text obtained from different sources for further processing, by cleaning, normalizing and splitting it into individual words or tokens.
Feature Extraction	Transform text using methods like Bag-of-Words, TF-IDF, Word2Vec and GloVe to extract features that you can use in machine learning models.
Modeling	A selection of different NLP tasks and how to build models that accomplish them.

Part 5: Voice User Interfaces

This module will help you get started in the exciting and fast-growing area of designing Voice User Interfaces! You'll learn how to build Conversational Agents for products and services more natural to interact with. You will also dive deeper into the core challenge of Speech Recognition, applying Recurrent Neural Networks to solve it.

Project: DNN Speech Recognizer

Build an Automatic Speech Recognizer using Deep Learning RNN's

Supporting Lessons

Lesson	Summary
Introduction to Speech Recognition	Dive deeper into the exciting field of Speech Recognition, including cutting edge deep learning technologies for Automatic Speech Recognition (ASR).



Udacity

Generated Tue Mar 31 17:11:25 PDT 2020