

Chapter 1 : Python: Real World Machine Learning - PDF Free Download - Fox eBook

This Learning Path will teach you Python machine learning for the real world. The machine learning techniques covered in this Learning Path are at the forefront of commercial practice. This Learning Path combines some of the best that Packt has to offer in one complete, curated package.

From the Foreword by Beau Cronin, 21 Inc. Real-World Machine Learning is a practical guide designed to teach working developers the art of ML project execution. Without overdosing you on academic theory and complex mathematics, it introduces the day-to-day practice of machine learning, preparing you to successfully build and deploy powerful ML systems. Listen to this book in liveAudio! Use the text to search and navigate the audio, or download the audio-only recording for portable offline listening. You can purchase or upgrade to liveAudio here or in liveBook. Table of Contents takes you straight to the book detailed table of contents Part 1: The machine-learning workflow 1. What is machine learning? How Machines Learn 1. Using Data to Make Decisions 1. The Machine Learning Approach 1. Five Advantages to Machine Learning 1. Learning a Model from Data 1. Evaluating Model Performance 1. Boosting Model Performance with Advanced Techniques 1. Data Pre-processing and Feature Engineering 1. Improving models continuously with online methods 1. Scaling Models with Data Volume and Velocity 1. Terms from this chapter 2. Which features should be included? Obtaining ground-truth for the target variable 2. How much training data is required? Is the training set representative enough? Pre-processing the data for modeling 2. Dealing with missing data 2.

Chapter 2 : Python Programming Tutorials

Real-World Machine Learning is a practical guide designed to teach working developers the art of ML project execution. Without overdosing you on academic theory and complex mathematics, it introduces the day-to-day practice of machine learning, preparing you to successfully build and deploy powerful ML systems.

A course where you will learn to implement cutting edge machine learning algorithms to solve real world problems. We have carefully selected the projects which will cover important aspect of Machine learning such as Supervised Learning, Unsupervised learning and Neural network with deep learning. It will be a course for serious developers but will be fun and engaging. You will learn step by step implementation and can be a professional ML developer after completing this course. What is Machine Learning? Machine Learning as a career Machine learning is the top emerging Job vertical among all technology jobs world wide. With new avenues of machine learning such as consumer behaviour analysis, marketing and sales forecasting, IT security, fraud detection , finance and office automation there is a big opportunity for trained Machine learning experts. Our course aims to provide expertise which you can directly transfer in your day to day Jobs. From credit card fraud detection to offering a more personalized marketing experience, machine learning can not only help but will also change the way we surf the web, drive our cars and maybe even shower. Machine Learning is currently one of the hottest trends on the market with companies big and small using learning algorithms to offer a more personalised user experience for their customers. As technology gets smarter and better, so will your algorithms. We have designed the perfect course for helping you get started with machine learning by helping you understand the key algorithms that are used when writing machine learning programs. You will learn how to write the codes and then see them in action. This course has been designed keeping in mind all types of learners. It includes something for everyone, from newbies to even advanced programmers. The course focuses on helping you actually learn how to code and write the algorithms, instead of simply focusing on the theoretical or the boring aspects of Machine Learning. The syllabus has been curated by experts and dedicates itself to bringing you the best approach approach to ML. The course has an example-based approach and will help you become more comfortable applying algorithms to real world problems. The course will focus on building 10 projects from scratch that will not only familiarize you with ML algorithms, but also will help you learn exactly how to use them. Tentative 10 Projects First 5 New Projects are the additional projects which will be Added to the Course as we reached our First stretch goal. After this project, the student should have the necessary foundation to begin building and deploying machine learning algorithms for natural language processing. Text Classification Text Classification Building on the foundation developed in the previous project, this tutorial will dive deeper into Natural Language Processing. Project 6 “ Stock Market Clustering Project This project will use a K-means clustering algorithm to identify related companies by finding correlations among stock market movements over a given time span. It will use the open source data from the Yahoo Finance Python module. Also, in this project you are going to do a PCA dimensionality reduction to plot the data on a 2D plot. K-means clustering algorithm Project 7 “ Breast Cancer Detection Breast Cancer Detection This second project will build a program to help detect breast cancer malignancies by using a support vector machine. You are also going to use the K-nearest neighbour algorithm to compare and contrast performance with the support vector machine. K-nearest neighbour algorithm Project 8 “ Board Game Review Board Game Review This project is actually going to be predicting board game reviews and in this project you are going to predict the average reviews on a board game based on characteristics such as difficulty, length, number of players. This will be accomplished by performing linear regression analysis. Linear Regression Analysis Project 9 “ Credit Card Fraud Detection In this project, you are going to do a credit card fraud detection and going to focus on anomaly detection by using probability densities. Anomaly Detection by Using Probability Densities Project 10 “ Diabetes Onset Detection This final project is going to be a diabetes onset detection with deep learning grid search. In this project, you will fine-tune a deep learning neural network by performing a grid search. The network will be used to detect the onset of diabetes based on patient data. Eduonix aims to identify and provide the best

learning and training environment. It identifies industry veterans and content creators around the globe and bring it to the global audience using number of intuitive platforms for easy and affordable access to quality content. Our Previous Successful Campaigns Our Previous Successful Campaigns Risks and challenges We have successfully delivered all our previous 8 projects on time, so we run a very small risk that we might be able to finish the course in the stipulated time. However, with multiple people working on this project day and night, we are confident that we will not pass the deadline. Questions about this project?

Chapter 3 : Python: Real World Machine Learning – CoderProg

In the first module, Python Machine Learning Cookbook, you will learn how to perform various machine learning tasks using a wide variety of machine learning algorithms to solve real-world problems and use Python to implement these algorithms.

Improve Health Care 1. This project is awesome for 3 main reasons: Which models are robust to missing data? Which models handle categorical features well? Tutorials Predicting wine quality with Scikit-Learn – Step-by-step tutorial for training a machine learning model R: There are some really fun datasets here, including PokemonGo spawn locations and Burritos in San Diego. Data for teams, games, scores, and players are all tracked and freely available online. There are plenty of fun machine learning projects for beginners. Sports is also an excellent domain for practicing data visualization and exploratory analysis. You can use these skills to help you decide which types of data to include in your analyses. More cluttered interface, but individual tables can be exported as CSV files. Predict Stock Prices The stock market is like candy-land for any data scientists who are even remotely interested in finance. You can find prices, fundamentals, global macroeconomic indicators, volatility indices, etc – the list goes on and on. Building trading models to practice machine learning is simple. Making them profitable is extremely difficult. Nothing here is financial advice, and we do not recommend trading real money. Quantopian – Quantitative finance community that offers a free platform for developing trading algorithm. Teach a Neural Network to Read Handwriting Neural networks and deep learning are two success stories in modern artificial intelligence. To start, we recommend with the first chapter in the tutorial below. It will teach you how to build a neural network from scratch that solves the MNIST challenge with high accuracy. The author also gives a very good explanation of the intuition behind neural networks. National Institute of Standards and Technology. It contains 70, labeled images of handwritten digits. In the year , Enron was one of the largest energy companies in America. Then, after being outed for fraud, it spiraled downward into bankruptcy within a year. Luckily for us, we have the Enron email database. It contains thousand emails between former Enron employees, mostly senior executives. In fact, data scientists have been using this dataset for education and research for years. Natural language processing – Analyze the body messages in conjunction with email metadata to classify emails based on their purposes. Write ML Algorithms from Scratch Writing machine learning algorithms from scratch is an excellent learning tool for two main reasons. Those packages are the fruits of years of development!

Chapter 4 : Python: Real World Machine Learning | PACKT Books

Python: Real World Machine Learning and millions of other books are available for Amazon Kindle. Learn more Enter your mobile number or email address below and we'll send you a link to download the free Kindle App.

Install numpy, matplotlib, pandas, sklearn and their dependencies Need help installing packages with pip? The objective of this course is to give you a wholistic understanding of machine learning, covering theory, application, and inner workings of supervised, unsupervised, and deep learning algorithms. For each major algorithm that we cover, we will discuss the high level intuitions of the algorithms and how they are logically meant to work. This should give you a complete understanding of exactly how the algorithms work, how they can be tweaked, what advantages are, and what their disadvantages are. In order to follow along with the series, I suggest you have at the very least a basic understanding of Python. If you do not, I suggest you at least follow the Python 3 Basics tutorial until the module installation with pip tutorial. Most of the machine learning algorithms are actually quite simple, since they need to be in order to scale to large datasets. Math involved is typically linear algebra, but I will do my best to still explain all of the math. Machine learning was defined in by Arthur Samuel as the "field of study that gives computers the ability to learn without being explicitly programmed. From what I have personally found, people outside the programming community mainly believe machine intelligence is hard-coded, completely unaware of the reality of the field. One of the largest challenges I had with machine learning was the abundance of material on the learning part. You can find formulas, charts, equations, and a bunch of theory on the topic of machine learning, but very little on the actual "machine" part, where you actually program the machine and run the algorithms on real data. This is mainly due to the history. In the 50s, machines were quite weak, and in very little supply, which remained very much the case for half a century. Machine Learning was relegated to being mainly theoretical and rarely actually employed. The "idea" of machine learning has come in and out of favor a few times through history, each time leaving people thinking it was merely a fad. Beyond this, there are ample resources out there to help you on your journey with machine learning, like this tutorial. You can just do a Google search on the topic and find more than enough information to keep you busy for a while. This is so much so to the point where we now have modules and APIs at our disposal, and you can engage in machine learning very easily without almost any knowledge at all of how it works. Machine learning is a lot like a car, you do not need to know much about how it works in order to get an incredible amount of utility from it. If you want to push the limits on performance and efficiency, however, you need to dig in under the hood, which is more how this course is geared. If you are just looking for a quick tutorial for employing machine learning on data, I already have a simple classification example tutorial and a simple clustering unsupervised machine learning example that you can check out. Machines are quite powerful, the one you are working on can probably do most of this series quickly. Data is also very plentiful lately. Make sure you have Python 3 installed, along with Pandas and Scikit-Learn.

Chapter 5 : Python books on Machine Learning and AI - calendrierdelascience.com

Use predictive modeling and apply it to real-world problems Understand how to perform market segmentation using unsupervised learning Apply your new found skills to solve real problems, through clearly-explained code for every technique and test.

Chapter 6 : Python: Real World Machine Learning [Book]

One of the largest challenges I had with machine learning was the abundance of material on the learning part. You can find formulas, charts, equations, and a bunch of theory on the topic of machine learning, but very little on the actual "machine" part, where you actually program the machine and run the algorithms on real data.

Chapter 7 : Manning | Real-World Machine Learning

About this course: This course dives into the basics of machine learning using an approachable, and well-known programming language, calendrierdelascience.com this course, we will be reviewing two main components: First, you will be learning about the purpose of Machine Learning and where it applies in the real world.

Chapter 8 : Preprocessing data using different techniques - Python: Real World Machine Learning [Book]

Download and install Python SciPy and get the most useful package for machine learning in Python. Load a dataset and understand it's structure using statistical summaries and data visualization. Create 6 machine learning models, pick the best and build confidence that the accuracy is reliable. If.

Chapter 9 : Real-World Machine Learning

From web development to data science, machine learning, and more, Python's real-world applications are limitless. Here are some projects that will assist you in finally putting your Python skills to good use.