

Hands On Transfer Learning With Python Implement

Getting the books **Hands On Transfer Learning With Python Implement** now is not type of inspiring means. You could not unaided going as soon as ebook hoard or library or borrowing from your associates to approach them. This is an enormously easy means to specifically get lead by on-line. This online broadcast Hands On Transfer Learning With Python Implement can be one of the options to accompany you with having other time.

It will not waste your time. say yes me, the e-book will utterly spread you extra business to read. Just invest little become old to way in this on-line declaration **Hands On Transfer Learning With Python Implement** as with ease as evaluation them wherever you are now.

Deep Learning for Computer Vision -
Jason Brownlee 2019-04-04
Step-by-step tutorials on deep learning
neural networks for computer vision in

python with Keras.
Deep Learning with Python - Francois
Chollet 2017-11-30
Summary Deep Learning with Python

introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language

and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View,

CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-

learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance

Python Deep Learning - Ivan Vasilev
2019-01-16

Learn advanced state-of-the-art deep learning techniques and their applications using popular Python libraries Key FeaturesBuild a strong foundation in neural networks and deep learning with Python librariesExplore advanced deep learning techniques and their applications across computer vision and NLP Learn how a computer can navigate in complex environments with reinforcement learningBook Description With the surge in artificial intelligence in applications catering to both business and consumer needs, deep learning is more important than ever for meeting current and future market

demands. With this book, you'll explore deep learning, and learn how to put machine learning to use in your projects. This second edition of Python Deep Learning will get you up to speed with deep learning, deep neural networks, and how to train them with high-performance algorithms and popular Python frameworks. You'll uncover different neural network architectures, such as convolutional networks, recurrent neural networks, long short-term memory (LSTM) networks, and capsule networks. You'll also learn how to solve problems in the fields of computer vision, natural language processing (NLP), and speech recognition. You'll study generative model approaches such as variational autoencoders and Generative Adversarial Networks (GANs) to generate images. As you delve into newly evolved areas of reinforcement learning, you'll gain an understanding of state-of-the-art algorithms that are the main

components behind popular games Go, Atari, and Dota. By the end of the book, you will be well-versed with the theory of deep learning along with its real-world applications. What you will learnGrasp the mathematical theory behind neural networks and deep learning processesInvestigate and resolve computer vision challenges using convolutional networks and capsule networksSolve generative tasks using variational autoencoders and Generative Adversarial NetworksImplement complex NLP tasks using recurrent networks (LSTM and GRU) and attention modelsExplore reinforcement learning and understand how agents behave in a complex environmentGet up to date with applications of deep learning in autonomous vehiclesWho this book is for This book is for data science practitioners, machine learning engineers, and those interested in deep learning who have a

basic foundation in machine learning and some Python programming experience. A background in mathematics and conceptual understanding of calculus and statistics will help you gain maximum benefit from this book.

Hands-On Automated Machine Learning

- Sibانjan Das 2018-04-26

Automate data and model pipelines for faster machine learning applications
Key Features
Build automated modules for different machine learning components
Understand each component of a machine learning pipeline in depth
Learn to use different open source AutoML and feature engineering platforms
Book Description
AutoML is designed to automate parts of Machine Learning. Readily available AutoML tools are making data science practitioners' work easy and are received well in the advanced analytics community. Automated Machine Learning covers the necessary

foundation needed to create automated machine learning modules and helps you get up to speed with them in the most practical way possible. In this book, you'll learn how to automate different tasks in the machine learning pipeline such as data preprocessing, feature selection, model training, model optimization, and much more. In addition to this, it demonstrates how you can use the available automation libraries, such as auto-sklearn and MLBox, and create and extend your own custom AutoML components for Machine Learning. By the end of this book, you will have a clearer understanding of the different aspects of automated Machine Learning, and you'll be able to incorporate automation tasks using practical datasets. You can leverage your learning from this book to implement Machine Learning in your projects and get a step closer to winning various machine learning competitions.

What you will learn Understand the fundamentals of Automated Machine Learning systems Explore auto-sklearn and MLBox for AutoML tasks Automate your preprocessing methods along with feature transformation Enhance feature selection and generation using the Python stack Assemble individual components of ML into a complete AutoML framework Demystify hyperparameter tuning to optimize your ML models Dive into Machine Learning concepts such as neural networks and autoencoders Understand the information costs and trade-offs associated with AutoML Who this book is for If you're a budding data scientist, data analyst, or Machine Learning enthusiast and are new to the concept of automated machine learning, this book is ideal for you. You'll also find this book useful if you're an ML engineer or data professional interested in developing quick machine learning pipelines for your projects. Prior exposure to

hands-on-transfer-learning-with-python-implement

Python programming will help you get the best out of this book.

Deep Learning with PyTorch - Luca Pietro Giovanni Antiga 2020-07-01

"We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document." —Soumith Chintala, co-creator of PyTorch Key Features Written by PyTorch's creator and key contributors Develop deep learning models in a familiar Pythonic way Use PyTorch to build an image classifier for cancer detection Diagnose problems with your neural network and improve training with data augmentation Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About The Book Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud

6/50

Downloaded from info.ucel.edu.ar on by @guest

detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you'll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks. What You Will Learn Understanding deep learning data structures such as tensors and neural networks Best practices for the

PyTorch Tensor API, loading data in Python, and visualizing results Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Methods for training networks with limited inputs Sifting through unreliable results to diagnose and fix problems in your neural network Improve your results with augmented data, better model architecture, and fine tuning This Book Is Written For For Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required. About The Authors Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning

and PyTorch speciality trainer and consultant based in Munich, Germany and a PyTorch core developer. Table of Contents
PART 1 - CORE PYTORCH 1 Introducing deep learning and the PyTorch Library 2 Pretrained networks 3 It starts with a tensor 4 Real-world data representation using tensors 5 The mechanics of learning 6 Using a neural network to fit the data 7 Telling birds from airplanes: Learning from images 8 Using convolutions to generalize PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER 9 Using PyTorch to fight cancer 10 Combining data sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production

Machine Learning for Algorithmic Trading - Stefan Jansen 2020-07-31

Leverage machine learning to design and back-test automated trading strategies for real-world markets using pandas, TA-Lib, scikit-learn, LightGBM, SpaCy, Gensim, TensorFlow 2, Zipline, backtrader, Alphalens, and pyfolio. Key FeaturesDesign, train, and evaluate machine learning algorithms that underpin automated trading strategiesCreate a research and strategy development process to apply predictive modeling to trading decisionsLeverage NLP and deep learning to extract tradeable signals from market and alternative dataBook Description The explosive growth of digital data has boosted the demand for expertise in trading strategies that use machine learning (ML). This revised and expanded second edition enables you to build and evaluate sophisticated supervised, unsupervised, and reinforcement learning

models. This book introduces end-to-end machine learning for the trading workflow, from the idea and feature engineering to model optimization, strategy design, and backtesting. It illustrates this by using examples ranging from linear models and tree-based ensembles to deep-learning techniques from cutting edge research. This edition shows how to work with market, fundamental, and alternative data, such as tick data, minute and daily bars, SEC filings, earnings call transcripts, financial news, or satellite images to generate tradeable signals. It illustrates how to engineer financial features or alpha factors that enable an ML model to predict returns from price data for US and international stocks and ETFs. It also shows how to assess the signal content of new features using Alphas and SHAP values and includes a new appendix with over one hundred alpha factor examples. By the end, you will be

proficient in translating ML model predictions into a trading strategy that operates at daily or intraday horizons, and in evaluating its performance. What you will learnLeverage market, fundamental, and alternative text and image dataResearch and evaluate alpha factors using statistics, Alphas, and SHAP valuesImplement machine learning techniques to solve investment and trading problemsBacktest and evaluate trading strategies based on machine learning using Zipline and BacktraderOptimize portfolio risk and performance analysis using pandas, NumPy, and pyfolioCreate a pairs trading strategy based on cointegration for US equities and ETFsTrain a gradient boosting model to predict intraday returns using AlgoSeek's high-quality trades and quotes dataWho this book is for If you are a data analyst, data scientist, Python developer, investment analyst, or portfolio manager interested in

getting hands-on machine learning knowledge for trading, this book is for you. This book is for you if you want to learn how to extract value from a diverse set of data sources using machine learning to design your own systematic trading strategies. Some understanding of Python and machine learning techniques is required.

Data Science and Data Analytics - Amit Kumar Tyagi 2021-09-22

Data science is a multi-disciplinary field that uses scientific methods, processes, algorithms, and systems to extract knowledge and insights from structured (labeled) and unstructured (unlabeled) data. It is the future of Artificial Intelligence (AI) and a necessity of the future to make things easier and more productive. In simple terms, data science is the discovery of data or uncovering hidden patterns (such as complex behaviors, trends, and inferences) from data. Moreover, Big Data

analytics/data analytics are the analysis mechanisms used in data science by data scientists. Several tools, such as Hadoop, R, etc., are used to analyze this large amount of data to predict valuable information and for decision-making. Note that structured data can be easily analyzed by efficient (available) business intelligence tools, while most of the data (80% of data by 2020) is in an unstructured form that requires advanced analytics tools. But while analyzing this data, we face several concerns, such as complexity, scalability, privacy leaks, and trust issues. Data science helps us to extract meaningful information or insights from unstructured or complex or large amounts of data (available or stored virtually in the cloud). Data Science and Data Analytics: Opportunities and Challenges covers all possible areas, applications with arising serious concerns, and challenges in this emerging field in

detail with a comparative analysis/taxonomy. FEATURES Gives the concept of data science, tools, and algorithms that exist for many useful applications Provides many challenges and opportunities in data science and data analytics that help researchers to identify research gaps or problems Identifies many areas and uses of data science in the smart era Applies data science to agriculture, healthcare, graph mining, education, security, etc. Academicians, data scientists, and stockbrokers from industry/business will find this book useful for designing optimal strategies to enhance their firm's productivity.

Hands-On Image Processing with Python - Sandipan Dey 2018-11-30

Explore the mathematical computations and algorithms for image processing using popular Python tools and frameworks. Key Features Practical coverage of every image

processing task with popular Python libraries Includes topics such as pseudo-coloring, noise smoothing, computing image descriptors Covers popular machine learning and deep learning techniques for complex image processing tasks Book Description Image processing plays an important role in our daily lives with various applications such as in social media (face detection), medical imaging (X-ray, CT-scan), security (fingerprint recognition) to robotics & space. This book will touch the core of image processing, from concepts to code using Python. The book will start from the classical image processing techniques and explore the evolution of image processing algorithms up to the recent advances in image processing or computer vision with deep learning. We will learn how to use image processing libraries such as PIL, scikit-image, and scipy ndimage in Python. This book will enable us to write code

snippets in Python 3 and quickly implement complex image processing algorithms such as image enhancement, filtering, segmentation, object detection, and classification. We will be able to use machine learning models using the scikit-learn library and later explore deep CNN, such as VGG-19 with Keras, and we will also use an end-to-end deep learning model called YOLO for object detection. We will also cover a few advanced problems, such as image inpainting, gradient blending, variational denoising, seam carving, quilting, and morphing. By the end of this book, we will have learned to implement various algorithms for efficient image processing. What you will learn

Perform basic data pre-processing tasks such as image denoising and spatial filtering in Python

Implement Fast Fourier Transform (FFT) and Frequency domain filters (e.g., Weiner) in Python

Do morphological image

processing and segment images with different algorithms

Learn techniques to extract features from images and match images

Write Python code to implement supervised / unsupervised machine learning algorithms for image processing

Use deep learning models for image classification, segmentation, object detection and style transfer

Who this book is for This book is for Computer Vision Engineers, and machine learning developers who are good with Python programming and want to explore details and complexities of image processing. No prior knowledge of the image processing techniques is expected.

Computational Science - ICCS 2022 - Derek Groen 2022-06-21

The four-volume set LNCS 13350, 13351, 13352, and 13353 constitutes the proceedings of the 22nd International Conference on Computational Science, ICCS 2022, held in London, UK, in June 2022.* The

total of 175 full papers and 78 short papers presented in this book set were carefully reviewed and selected from 474 submissions. 169 full and 36 short papers were accepted to the main track; 120 full and 42 short papers were accepted to the workshops/ thematic tracks. *The conference was held in a hybrid format

Practical Machine Learning with Python

- Dipanjan Sarkar 2017-12-20

Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning

systems and projects successfully. Practical Machine Learning with Python follows a structured and comprehensive three-tiered approach packed with hands-on examples and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard machine learning pipelines, with an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize and visualize data in its various forms. Feature engineering and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and

deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies, music, marketing, computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. Practical Machine Learning with Python will empower you to start solving your own problems with machine learning today! What You'll Learn Execute end-to-end machine learning projects and systems Implement hands-on examples with industry standard, open source, robust machine learning tools and frameworks Review case studies depicting applications of machine learning and deep learning on diverse domains and industries Apply a wide

hands-on-transfer-learning-with-python-implement

range of machine learning models including regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including CNNs, RNNs, LSTMs and transfer learning. Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow -

Aurélien Géron 2019-09-05

Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—Scikit-Learn and TensorFlow—author Aurélien Géron helps

14/50

Downloaded from info.ucl.edu.ar on by @guest

you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets

Deep Learning with PyTorch Lightning -
Kunal Sawarkar 2022-02-04

Build, train, deploy, and scale deep learning models quickly and accurately, improving your productivity using the lightweight PyTorch Wrapper Key Features: Become well-versed with PyTorch Lightning architecture and learn how it can be implemented in various industry domains Speed up your research using PyTorch Lightning by creating new loss functions, networks, and architectures Train and build new algorithms for massive data using distributed training Book Description: PyTorch Lightning lets researchers build their own Deep Learning (DL) models without having to worry about the boilerplate. With the help of this book, you'll be able to maximize productivity for DL projects while ensuring full flexibility from model formulation through to implementation. You'll take a hands-on approach to implementing PyTorch Lightning models to get up to speed in no

time. You'll start by learning how to configure PyTorch Lightning on a cloud platform, understand the architectural components, and explore how they are configured to build various industry solutions. Next, you'll build a network and application from scratch and see how you can expand it based on your specific needs, beyond what the framework can provide. The book also demonstrates how to implement out-of-box capabilities to build and train Self-Supervised Learning, semi-supervised learning, and time series models using PyTorch Lightning. As you advance, you'll discover how generative adversarial networks (GANs) work. Finally, you'll work with deployment-ready applications, focusing on faster performance and scaling, model scoring on massive volumes of data, and model debugging. By the end of this PyTorch book, you'll have developed the knowledge and skills necessary to build and

deploy your own scalable DL applications using PyTorch Lightning. What You Will Learn: Customize models that are built for different datasets, model architectures, and optimizers Understand how a variety of Deep Learning models from image recognition and time series to GANs, semi-supervised and self-supervised models can be built Use out-of-the-box model architectures and pre-trained models using transfer learning Run and tune DL models in a multi-GPU environment using mixed-mode precisions Explore techniques for model scoring on massive workloads Discover troubleshooting techniques while debugging DL models Who this book is for: This deep learning book is for citizen data scientists and expert data scientists transitioning from other frameworks to PyTorch Lightning. This book will also be useful for deep learning researchers who are just getting started with coding for deep learning models using

PyTorch Lightning. Working knowledge of Python programming and an intermediate-level understanding of statistics and deep learning fundamentals is expected.

Deep Learning for Coders with fastai and PyTorch - Jeremy Howard 2020-06-29

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of

the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

Practical Machine Learning with Python

- Dipanjan Sarkar 2017-12-22

Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science

of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning systems and projects successfully. Practical Machine Learning with Python follows a structured and comprehensive three-tiered approach packed with hands-on examples and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard machine learning pipelines, with an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize

and visualize data in its various forms. Feature engineering and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies, music, marketing, computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. Practical Machine Learning with Python will empower you to start solving your own problems with machine learning today! What You'll Learn Execute end-to-end machine learning projects and systems Implement hands-on examples with

industry standard, open source, robust machine learning tools and frameworks Review case studies depicting applications of machine learning and deep learning on diverse domains and industries Apply a wide range of machine learning models including regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including CNNs, RNNs, LSTMs and transfer learning. Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students
Hands-On Convolutional Neural Networks with TensorFlow - Iffat Zafar 2018-08-28 Learn how to apply TensorFlow to a wide range of deep learning and Machine Learning problems with this practical guide on training CNNs for image classification, image recognition, object detection and many computer vision challenges. Key Features Learn the fundamentals of

hands-on-transfer-learning-with-python-implement

Convolutional Neural Networks Harness Python and Tensorflow to train CNNs Build scalable deep learning models that can process millions of items Book Description Convolutional Neural Networks (CNN) are one of the most popular architectures used in computer vision apps. This book is an introduction to CNNs through solving real-world problems in deep learning while teaching you their implementation in popular Python library - TensorFlow. By the end of the book, you will be training CNNs in no time! We start with an overview of popular machine learning and deep learning models, and then get you set up with a TensorFlow development environment. This environment is the basis for implementing and training deep learning models in later chapters. Then, you will use Convolutional Neural Networks to work on problems such as image classification, object detection, and semantic segmentation. After that, you

19/50

Downloaded from info.ucel.edu.ar on by @guest

will use transfer learning to see how these models can solve other deep learning problems. You will also get a taste of implementing generative models such as autoencoders and generative adversarial networks. Later on, you will see useful tips on machine learning best practices and troubleshooting. Finally, you will learn how to apply your models on large datasets of millions of images. What you will learn Train machine learning models with TensorFlow Create systems that can evolve and scale during their life cycle Use CNNs in image recognition and classification Use TensorFlow for building deep learning models Train popular deep learning models Fine-tune a neural network to improve the quality of results with transfer learning Build TensorFlow models that can scale to large datasets and systems Who this book is for This book is for Software Engineers, Data Scientists, or Machine Learning practitioners

who want to use CNNs for solving real-world problems. Knowledge of basic machine learning concepts, linear algebra and Python will help.

Modern Computer Vision with PyTorch -

V Kishore Ayyadevara 2020-11-27

Get to grips with deep learning techniques for building image processing applications using PyTorch with the help of code notebooks and test questions Key FeaturesImplement solutions to 50 real-world computer vision applications using PyTorchUnderstand the theory and working mechanisms of neural network architectures and their implementationDiscover best practices using a custom library created especially for this bookBook Description Deep learning is the driving force behind many recent advances in various computer vision (CV) applications. This book takes a hands-on approach to help you to solve over 50 CV problems using PyTorch1.x on real-

world datasets. You'll start by building a neural network (NN) from scratch using NumPy and PyTorch and discover best practices for tweaking its hyperparameters. You'll then perform image classification using convolutional neural networks and transfer learning and understand how they work. As you progress, you'll implement multiple use cases of 2D and 3D multi-object detection, segmentation, human-pose-estimation by learning about the R-CNN family, SSD, YOLO, U-Net architectures, and the Detectron2 platform. The book will also guide you in performing facial expression swapping, generating new faces, and manipulating facial expressions as you explore autoencoders and modern generative adversarial networks. You'll learn how to combine CV with NLP techniques, such as LSTM and transformer, and RL techniques, such as Deep Q-learning, to implement OCR, image captioning, object

detection, and a self-driving car agent. Finally, you'll move your NN model to production on the AWS Cloud. By the end of this book, you'll be able to leverage modern NN architectures to solve over 50 real-world CV problems confidently. What you will learnTrain a NN from scratch with NumPy and PyTorchImplement 2D and 3D multi-object detection and segmentationGenerate digits and DeepFakes with autoencoders and advanced GANsManipulate images using CycleGAN, Pix2PixGAN, StyleGAN2, and SRGANCombine CV with NLP to perform OCR, image captioning, and object detectionCombine CV with reinforcement learning to build agents that play pong and self-drive a carDeploy a deep learning model on the AWS server using FastAPI and DockerImplement over 35 NN architectures and common OpenCV utilitiesWho this book is for This book is for beginners to PyTorch and intermediate-level machine learning

practitioners who are looking to get well-versed with computer vision techniques using deep learning and PyTorch. If you are just getting started with neural networks, you'll find the use cases accompanied by notebooks in GitHub present in this book useful. Basic knowledge of the Python programming language and machine learning is all you need to get started with this book.

[The TensorFlow Workshop](#) - Matthew Moocarme 2021-12-15

Get started with TensorFlow fundamentals to build and train deep learning models with real-world data, practical exercises, and challenging activities

Key Features Understand the fundamentals of tensors, neural networks, and deep learning Discover how to implement and fine-tune deep learning models for real-world datasets Build your experience and confidence with hands-on exercises and

activities

Book Description Getting to grips with tensors, deep learning, and neural networks can be intimidating and confusing for anyone, no matter their experience level. The breadth of information out there, often written at a very high level and aimed at advanced practitioners, can make getting started even more challenging. If this sounds familiar to you, The TensorFlow Workshop is here to help. Combining clear explanations, realistic examples, and plenty of hands-on practice, it'll quickly get you up and running. You'll start off with the basics – learning how to load data into TensorFlow, perform tensor operations, and utilize common optimizers and activation functions. As you progress, you'll experiment with different TensorFlow development tools, including TensorBoard, TensorFlow Hub, and Google Colab, before moving on to solve regression and classification problems with sequential

models. Building on this solid foundation, you'll learn how to tune models and work with different types of neural network, getting hands-on with real-world deep learning applications such as text encoding, temperature forecasting, image augmentation, and audio processing. By the end of this deep learning book, you'll have the skills, knowledge, and confidence to tackle your own ambitious deep learning projects with TensorFlow. What you will learn

- Get to grips with TensorFlow's mathematical operations
- Pre-process a wide variety of tabular, sequential, and image data
- Understand the purpose and usage of different deep learning layers
- Perform hyperparameter-tuning to prevent overfitting of training data
- Use pre-trained models to speed up the development of learning models
- Generate new data based on existing patterns using generative models

Who this book is for This TensorFlow

book is for anyone who wants to develop their understanding of deep learning and get started building neural networks with TensorFlow. Basic knowledge of Python programming and its libraries, as well as a general understanding of the fundamentals of data science and machine learning, will help you grasp the topics covered in this book more easily.

Practical Convolutional Neural Networks - Mohit Sewak 2018-02-27

This book helps you master CNN, from the basics to the most advanced concepts in CNN such as GANs, instance classification and attention mechanism for vision models and more. You will implement advanced CNN models using complex image and video datasets. By the end of the book you will learn CNN's best practices to implement smart ConvNet ...

Pro Machine Learning Algorithms - V Kishore Ayyadevara 2018-06-30

Bridge the gap between a high-level understanding of how an algorithm works and knowing the nuts and bolts to tune your models better. This book will give you the confidence and skills when developing all the major machine learning models. In *Pro Machine Learning Algorithms*, you will first develop the algorithm in Excel so that you get a practical understanding of all the levers that can be tuned in a model, before implementing the models in Python/R. You will cover all the major algorithms: supervised and unsupervised learning, which include linear/logistic regression; k-means clustering; PCA; recommender system; decision tree; random forest; GBM; and neural networks. You will also be exposed to the latest in deep learning through CNNs, RNNs, and word2vec for text mining. You will be learning not only the algorithms, but also the concepts of feature engineering to maximize the performance of

a model. You will see the theory along with case studies, such as sentiment classification, fraud detection, recommender systems, and image recognition, so that you get the best of both theory and practice for the vast majority of the machine learning algorithms used in industry. Along with learning the algorithms, you will also be exposed to running machine-learning models on all the major cloud service providers. You are expected to have minimal knowledge of statistics/software programming and by the end of this book you should be able to work on a machine learning project with confidence. What You Will Learn Get an in-depth understanding of all the major machine learning and deep learning algorithms Fully appreciate the pitfalls to avoid while building models Implement machine learning algorithms in the cloud Follow a hands-on approach through case studies for each algorithm

Gain the tricks of ensemble learning to build more accurate models Discover the basics of programming in R/Python and the Keras framework for deep learning Who This Book Is For Business analysts/ IT professionals who want to transition into data science roles. Data scientists who want to solidify their knowledge in machine learning.

Python Deep Learning - Valentino Zocca
2017-04-28

Take your machine learning skills to the next level by mastering Deep Learning concepts and algorithms using Python. About This Book Explore and create intelligent systems using cutting-edge deep learning techniques Implement deep learning algorithms and work with revolutionary libraries in Python Get real-world examples and easy-to-follow tutorials on Theano, TensorFlow, H2O and more Who This Book Is For This book is for Data Science practitioners as well as aspirants

who have a basic foundational understanding of Machine Learning concepts and some programming experience with Python. A mathematical background with a conceptual understanding of calculus and statistics is also desired. What You Will Learn Get a practical deep dive into deep learning algorithms Explore deep learning further with Theano, Caffe, Keras, and TensorFlow Learn about two of the most powerful techniques at the core of many practical deep learning implementations: Auto-Encoders and Restricted Boltzmann Machines Dive into Deep Belief Nets and Deep Neural Networks Discover more deep learning algorithms with Dropout and Convolutional Neural Networks Get to know device strategies so you can use deep learning algorithms and libraries in the real world In Detail With an increasing interest in AI around the world, deep learning has

attracted a great deal of public attention. Every day, deep learning algorithms are used broadly across different industries. The book will give you all the practical information available on the subject, including the best practices, using real-world use cases. You will learn to recognize and extract information to increase predictive accuracy and optimize results. Starting with a quick recap of important machine learning concepts, the book will delve straight into deep learning principles using Sci-kit learn. Moving ahead, you will learn to use the latest open source libraries such as Theano, Keras, Google's TensorFlow, and H2O. Use this guide to uncover the difficulties of pattern recognition, scaling data with greater accuracy and discussing deep learning algorithms and techniques. Whether you want to dive deeper into Deep Learning, or want to investigate how to get more out of

this powerful technology, you'll find everything inside. Style and approach Python Machine Learning by example follows practical hands on approach. It walks you through the key elements of Python and its powerful machine learning libraries with the help of real world projects.

[Programming with TensorFlow](#) - Kolla Bhanu Prakash 2021-01-22

This practical book provides an end-to-end guide to TensorFlow, the leading open source software library that helps you build and train neural networks for deep learning, Natural Language Processing (NLP), speech recognition, and general predictive analytics. The book provides a hands-on approach to TensorFlow fundamentals for a broad technical audience—from data scientists and engineers to students and researchers. The authors begin by working through some basic examples in TensorFlow before diving deeper into topics such as

CNN, RNN, LSTM, and GNN. The book is written for those who want to build powerful, robust, and accurate predictive models with the power of TensorFlow, combined with other open source Python libraries. The authors demonstrate TensorFlow projects on Single Board Computers (SBCs).

Supervised Machine Learning with Python - Taylor Smith 2019-05-27

Teach your machine to think for itself! Key FeaturesDelve into supervised learning and grasp how a machine learns from dataImplement popular machine learning algorithms from scratch, developing a deep understanding along the wayExplore some of the most popular scientific and mathematical libraries in the Python languageBook Description Supervised machine learning is used in a wide range of sectors (such as finance, online advertising, and analytics) because it allows you to train

your system to make pricing predictions, campaign adjustments, customer recommendations, and much more while the system self-adjusts and makes decisions on its own. As a result, it's crucial to know how a machine "learns" under the hood. This book will guide you through the implementation and nuances of many popular supervised machine learning algorithms while facilitating a deep understanding along the way. You'll embark on this journey with a quick overview and see how supervised machine learning differs from unsupervised learning. Next, we explore parametric models such as linear and logistic regression, non-parametric methods such as decision trees, and various clustering techniques to facilitate decision-making and predictions. As we proceed, you'll work hands-on with recommender systems, which are widely used by online companies to increase user interaction and

enrich shopping potential. Finally, you'll wrap up with a brief foray into neural networks and transfer learning. By the end of this book, you'll be equipped with hands-on techniques and will have gained the practical know-how you need to quickly and powerfully apply algorithms to new problems. What you will learn

Crack how a machine learns a concept and generalize its understanding to new data

Uncover the fundamental differences between parametric and non-parametric models

Implement and grok several well-known supervised learning algorithms from scratch

Work with models in domains such as ecommerce and marketing

Expand your expertise and use various algorithms such as regression, decision trees, and clustering

Build your own models capable of making predictions

Delve into the most popular approaches in deep learning such as transfer learning and neural networks

Who

this book is for This book is for aspiring machine learning developers who want to get started with supervised learning. Intermediate knowledge of Python programming—and some fundamental knowledge of supervised learning—are expected.

Hands-On Meta Learning with Python -

Sudharsan Ravichandiran 2018-12-31

Explore a diverse set of meta-learning algorithms and techniques to enable human-like cognition for your machine learning models using various Python frameworks

Key Features

Understand the foundations of meta learning algorithms

Explore practical examples to explore various one-shot learning algorithms with its applications in TensorFlow

Master state of the art meta learning algorithms like MAML, reptile, meta

SGD

Book Description

Meta learning is an exciting research trend in machine learning, which

enables a model to understand the learning process. Unlike other ML paradigms, with meta learning you can learn from small datasets faster. Hands-On Meta Learning with Python starts by explaining the fundamentals of meta learning and helps you understand the concept of learning to learn. You will delve into various one-shot learning algorithms, like siamese, prototypical, relation and memory-augmented networks by implementing them in TensorFlow and Keras. As you make your way through the book, you will dive into state-of-the-art meta learning algorithms such as MAML, Reptile, and CAML. You will then explore how to learn quickly with Meta-SGD and discover how you can perform unsupervised learning using meta learning with CACTUs. In the concluding chapters, you will work through recent trends in meta learning such as adversarial meta learning, task agnostic meta learning, and meta

imitation learning. By the end of this book, you will be familiar with state-of-the-art meta learning algorithms and able to enable human-like cognition for your machine learning models. What you will learn

Understand the basics of meta learning methods, algorithms, and types

Build voice and face recognition models using a siamese network

Learn the prototypical network along with its variants

Build relation networks and matching networks from scratch

Implement MAML and Reptile algorithms from scratch in Python

Work through imitation learning and adversarial meta learning

Explore task agnostic meta learning and deep meta learning

Who this book is for

Hands-On Meta Learning with Python is for machine learning enthusiasts, AI researchers, and data scientists who want to explore meta learning as an advanced approach for training machine learning models. Working knowledge of machine

learning concepts and Python programming is necessary.

Deep Learning By Example - Ahmed Menshawy 2018-02-28

Grasp the fundamental concepts of deep learning using Tensorflow in a hands-on manner Key Features Get a first-hand experience of the deep learning concepts and techniques with this easy-to-follow guide Train different types of neural networks using Tensorflow for real-world problems in language processing, computer vision, transfer learning, and more Designed for those who believe in the concept of 'learn by doing', this book is a perfect blend of theory and code examples Book Description Deep learning is a popular subset of machine learning, and it allows you to build complex models that are faster and give more accurate predictions. This book is your companion to take your first steps into the world of deep learning, with

hands-on examples to boost your understanding of the topic. This book starts with a quick overview of the essential concepts of data science and machine learning which are required to get started with deep learning. It introduces you to Tensorflow, the most widely used machine learning library for training deep learning models. You will then work on your first deep learning problem by training a deep feed-forward neural network for digit classification, and move on to tackle other real-world problems in computer vision, language processing, sentiment analysis, and more. Advanced deep learning models such as generative adversarial networks and their applications are also covered in this book. By the end of this book, you will have a solid understanding of all the essential concepts in deep learning. With the help of the examples and code provided in this book, you will be equipped to train your own

deep learning models with more confidence. What you will learn Understand the fundamentals of deep learning and how it is different from machine learning Get familiarized with Tensorflow, one of the most popular libraries for advanced machine learning Increase the predictive power of your model using feature engineering Understand the basics of deep learning by solving a digit classification problem of MNIST Demonstrate face generation based on the CelebA database, a promising application of generative models Apply deep learning to other domains like language modeling, sentiment analysis, and machine translation Who this book is for This book targets data scientists and machine learning developers who wish to get started with deep learning. If you know what deep learning is but are not quite sure of how to use it, this book will help you as well. An understanding of statistics and data science

concepts is required. Some familiarity with Python programming will also be beneficial. **Learn TensorFlow 2.0** - Pramod Singh 2019-12-17

Learn how to use TensorFlow 2.0 to build machine learning and deep learning models with complete examples. The book begins with introducing TensorFlow 2.0 framework and the major changes from its last release. Next, it focuses on building Supervised Machine Learning models using TensorFlow 2.0. It also demonstrates how to build models using customer estimators. Further, it explains how to use TensorFlow 2.0 API to build machine learning and deep learning models for image classification using the standard as well as custom parameters. You'll review sequence predictions, saving, serving, deploying, and standardized datasets, and then deploy these models to production. All the code presented in the book will be available in the form of

executable scripts at Github which allows you to try out the examples and extend them in interesting ways. What You'll Learn Review the new features of TensorFlow 2.0 Use TensorFlow 2.0 to build machine learning and deep learning models Perform sequence predictions using TensorFlow 2.0 Deploy TensorFlow 2.0 models with practical examples Who This Book Is For Data scientists, machine and deep learning engineers.

Hands-on Scikit-Learn for Machine Learning Applications - David Paper 2019-11-16 Aspiring data science professionals can learn the Scikit-Learn library along with the fundamentals of machine learning with this book. The book combines the Anaconda Python distribution with the popular Scikit-Learn library to demonstrate a wide range of supervised and unsupervised machine learning algorithms. Care is taken to walk you through the principles of machine

learning through clear examples written in Python that you can try out and experiment with at home on your own machine. All applied math and programming skills required to master the content are covered in this book. In-depth knowledge of object-oriented programming is not required as working and complete examples are provided and explained. Coding examples are in-depth and complex when necessary. They are also concise, accurate, and complete, and complement the machine learning concepts introduced. Working the examples helps to build the skills necessary to understand and apply complex machine learning algorithms. Hands-on Scikit-Learn for Machine Learning Applications is an excellent starting point for those pursuing a career in machine learning. Students of this book will learn the fundamentals that are a prerequisite to competency. Readers will be exposed to the Anaconda distribution of

Python that is designed specifically for data science professionals, and will build skills in the popular Scikit-Learn library that underlies many machine learning applications in the world of Python. What You'll Learn Work with simple and complex datasets common to Scikit-Learn Manipulate data into vectors and matrices for algorithmic processing Become familiar with the Anaconda distribution used in data science Apply machine learning with Classifiers, Regressors, and Dimensionality Reduction Tune algorithms and find the best algorithms for each dataset Load data from and save to CSV, JSON, Numpy, and Pandas formats Who This Book Is For The aspiring data scientist yearning to break into machine learning through mastering the underlying fundamentals that are sometimes skipped over in the rush to be productive. Some knowledge of object-oriented programming and very basic

applied linear algebra will make learning easier, although anyone can benefit from this book.

Programming PyTorch for Deep Learning - Ian Pointer 2019-09-20

Deep learning is changing everything. This machine-learning method has already surpassed traditional computer vision techniques, and the same is happening with NLP. If you're looking to bring deep learning into your domain, this practical book will bring you up to speed on key concepts using Facebook's PyTorch framework. Once author Ian Pointer helps you set up PyTorch on a cloud-based environment, you'll learn how use the framework to create neural architectures for performing operations on images, sound, text, and other types of data. By the end of the book, you'll be able to create neural networks and train them on multiple types of data. Learn how to deploy deep learning models to production Explore

PyTorch use cases in companies other than Facebook Learn how to apply transfer learning to images Apply cutting-edge NLP techniques using a model trained on Wikipedia

Hands-On Deep Learning for Games -

Micheal Lanham 2019-03-30

Understand the core concepts of deep learning and deep reinforcement learning by applying them to develop games Key Features Apply the power of deep learning to complex reasoning tasks by building a Game AI Exploit the most recent developments in machine learning and AI for building smart games Implement deep learning models and neural networks with Python Book Description The number of applications of deep learning and neural networks has multiplied in the last couple of years. Neural nets has enabled significant breakthroughs in everything from computer vision, voice generation, voice recognition

and self-driving cars. Game development is also a key area where these techniques are being applied. This book will give an in depth view of the potential of deep learning and neural networks in game development. We will take a look at the foundations of multi-layer perceptron's to using convolutional and recurrent networks. In applications from GANs that create music or textures to self-driving cars and chatbots. Then we introduce deep reinforcement learning through the multi-armed bandit problem and other OpenAI Gym environments. As we progress through the book we will gain insights about DRL techniques such as Motivated Reinforcement Learning with Curiosity and Curriculum Learning. We also take a closer look at deep reinforcement learning and in particular the Unity ML-Agents toolkit. By the end of the book, we will look at how to apply DRL and the ML-Agents toolkit to

enhance, test and automate your games or simulations. Finally, we will cover your possible next steps and possible areas for future learning. What you will learn Learn the foundations of neural networks and deep learning. Use advanced neural network architectures in applications to create music, textures, self driving cars and chatbots. Understand the basics of reinforcement and DRL and how to apply it to solve a variety of problems. Working with Unity ML-Agents toolkit and how to install, setup and run the kit. Understand core concepts of DRL and the differences between discrete and continuous action environments. Use several advanced forms of learning in various scenarios from developing agents to testing games. Who this book is for This book is for game developers who wish to create highly interactive games by leveraging the power of machine and deep learning. No prior

knowledge of machine learning, deep learning or neural networks is required this book will teach those concepts from scratch. A good understanding of Python is required. Hands-On Machine Learning for Algorithmic Trading - Stefan Jansen 2018-12-31 Explore effective trading strategies in real-world markets using NumPy, spaCy, pandas, scikit-learn, and Keras Key Features Implement machine learning algorithms to build, train, and validate algorithmic models Create your own algorithmic design process to apply probabilistic machine learning approaches to trading decisions Develop neural networks for algorithmic trading to perform time series forecasting and smart analytics Book Description The explosive growth of digital data has boosted the demand for expertise in trading strategies that use machine learning (ML). This book enables you to use a broad range of supervised and

unsupervised algorithms to extract signals from a wide variety of data sources and create powerful investment strategies. This book shows how to access market, fundamental, and alternative data via API or web scraping and offers a framework to evaluate alternative data. You'll practice the ML workflow from model design, loss metric definition, and parameter tuning to performance evaluation in a time series context. You will understand ML algorithms such as Bayesian and ensemble methods and manifold learning, and will know how to train and tune these models using pandas, statsmodels, sklearn, PyMC3, xgboost, lightgbm, and catboost. This book also teaches you how to extract features from text data using spaCy, classify news and assign sentiment scores, and to use gensim to model topics and learn word embeddings from financial reports. You will also build and evaluate neural networks, including

RNNs and CNNs, using Keras and PyTorch to exploit unstructured data for sophisticated strategies. Finally, you will apply transfer learning to satellite images to predict economic activity and use reinforcement learning to build agents that learn to trade in the OpenAI Gym. What you will learnImplement machine learning techniques to solve investment and trading problemsLeverage market, fundamental, and alternative data to research alpha factorsDesign and fine-tune supervised, unsupervised, and reinforcement learning modelsOptimize portfolio risk and performance using pandas, NumPy, and scikit-learnIntegrate machine learning models into a live trading strategy on QuantopianEvaluate strategies using reliable backtesting methodologies for time seriesDesign and evaluate deep neural networks using Keras, PyTorch, and TensorFlowWork with reinforcement learning

for trading strategies in the OpenAI GymWho this book is for Hands-On Machine Learning for Algorithmic Trading is for data analysts, data scientists, and Python developers, as well as investment analysts and portfolio managers working within the finance and investment industry. If you want to perform efficient algorithmic trading by developing smart investigating strategies using machine learning algorithms, this is the book for you. Some understanding of Python and machine learning techniques is mandatory.

Hands-On One-shot Learning with Python - Shruti Jadon 2020-04-10

Get to grips with building powerful deep learning models using PyTorch and scikit-learn Key FeaturesLearn how you can speed up the deep learning process with one-shot learningUse Python and PyTorch to build state-of-the-art one-shot learning modelsExplore architectures such as

Siamese networks, memory-augmented neural networks, model-agnostic meta-learning, and discriminative k-shot learningBook Description One-shot learning has been an active field of research for scientists trying to develop a cognitive machine that mimics human learning. With this book, you'll explore key approaches to one-shot learning, such as metrics-based, model-based, and optimization-based techniques, all with the help of practical examples. Hands-On One-shot Learning with Python will guide you through the exploration and design of deep learning models that can obtain information about an object from one or just a few training samples. The book begins with an overview of deep learning and one-shot learning and then introduces you to the different methods you can use to achieve it, such as deep learning architectures and probabilistic models. Once you've got to grips with the

core principles, you'll explore real-world examples and implementations of one-shot learning using PyTorch 1.x on datasets such as Omniglot and MiniImageNet. Finally, you'll explore generative modeling-based methods and discover the key considerations for building systems that exhibit human-level intelligence. By the end of this book, you'll be well-versed with the different one- and few-shot learning methods and be able to use them to build your own deep learning models. What you will learn

Get to grips with the fundamental concepts of one- and few-shot learning

Work with different deep learning architectures for one-shot learning

Understand when to use one-shot and transfer learning, respectively

Study the Bayesian network approach for one-shot learning

Implement one-shot learning approaches based on metrics, models, and optimization in PyTorch

Discover different optimization

algorithms that help to improve accuracy even with smaller volumes of data

Explore various one-shot learning architectures based on classification and regression

Who this book is for

If you're an AI researcher or a machine learning or deep learning expert looking to explore one-shot learning, this book is for you. It will help you get started with implementing various one-shot techniques to train models faster. Some Python programming experience is necessary to understand the concepts covered in this book.

Better Deep Learning - Jason Brownlee
2018-12-13

Deep learning neural networks have become easy to define and fit, but are still hard to configure. Discover exactly how to improve the performance of deep learning neural network models on your predictive modeling projects. With clear explanations, standard Python libraries, and step-by-step tutorial

lessons, you'll discover how to better train your models, reduce overfitting, and make more accurate predictions.

Transfer Learning - Qiang Yang

2020-01-31

Transfer learning deals with how systems can quickly adapt themselves to new situations, tasks and environments. It gives machine learning systems the ability to leverage auxiliary data and models to help solve target problems when there is only a small amount of data available. This makes such systems more reliable and robust, keeping the machine learning model faced with unforeseeable changes from deviating too much from expected performance. At an enterprise level, transfer learning allows knowledge to be reused so experience gained once can be repeatedly applied to the real world. For example, a pre-trained model that takes account of user privacy can be downloaded and adapted at the edge

of a computer network. This self-contained, comprehensive reference text describes the standard algorithms and demonstrates how these are used in different transfer learning paradigms. It offers a solid grounding for newcomers as well as new insights for seasoned researchers and developers.

Python Deep Learning Cookbook - Indra den Bakker 2017-10-27

Solve different problems in modelling deep neural networks using Python, Tensorflow, and Keras with this practical guide About This Book Practical recipes on training different neural network models and tuning them for optimal performance Use Python frameworks like TensorFlow, Caffe, Keras, Theano for Natural Language Processing, Computer Vision, and more A hands-on guide covering the common as well as the not so common problems in deep learning using Python Who This Book Is For This book is intended for machine learning

professionals who are looking to use deep learning algorithms to create real-world applications using Python. Thorough understanding of the machine learning concepts and Python libraries such as NumPy, SciPy and scikit-learn is expected. Additionally, basic knowledge in linear algebra and calculus is desired. What You Will Learn Implement different neural network models in Python Select the best Python framework for deep learning such as PyTorch, Tensorflow, MXNet and Keras Apply tips and tricks related to neural networks internals, to boost learning performances Consolidate machine learning principles and apply them in the deep learning field Reuse and adapt Python code snippets to everyday problems Evaluate the cost/benefits and performance implication of each discussed solution In Detail Deep Learning is revolutionizing a wide range of industries. For many applications, deep learning has

proven to outperform humans by making faster and more accurate predictions. This book provides a top-down and bottom-up approach to demonstrate deep learning solutions to real-world problems in different areas. These applications include Computer Vision, Natural Language Processing, Time Series, and Robotics. The Python Deep Learning Cookbook presents technical solutions to the issues presented, along with a detailed explanation of the solutions. Furthermore, a discussion on corresponding pros and cons of implementing the proposed solution using one of the popular frameworks like TensorFlow, PyTorch, Keras and CNTK is provided. The book includes recipes that are related to the basic concepts of neural networks. All techniques, as well as classical networks topologies. The main purpose of this book is to provide Python programmers a detailed list of recipes to apply deep learning to common

and not-so-common scenarios. Style and approach Unique blend of independent recipes arranged in the most logical manner

[Hands-On Transfer Learning with Python](#) -

Dipanjan Sarkar 2018-08-31

Deep learning simplified by taking supervised, unsupervised, and reinforcement learning to the next level using the Python ecosystem Key Features Build deep learning models with transfer learning principles in Python implement transfer learning to solve real-world research problems Perform complex operations such as image captioning neural style transfer Book Description Transfer learning is a machine learning (ML) technique where knowledge gained during training a set of problems can be used to solve other similar problems. The purpose of this book is two-fold; firstly, we focus on detailed coverage of deep learning (DL) and transfer learning, comparing and contrasting

the two with easy-to-follow concepts and examples. The second area of focus is real-world examples and research problems using TensorFlow, Keras, and the Python ecosystem with hands-on examples. The book starts with the key essential concepts of ML and DL, followed by depiction and coverage of important DL architectures such as convolutional neural networks (CNNs), deep neural networks (DNNs), recurrent neural networks (RNNs), long short-term memory (LSTM), and capsule networks. Our focus then shifts to transfer learning concepts, such as model freezing, fine-tuning, pre-trained models including VGG, inception, ResNet, and how these systems perform better than DL models with practical examples. In the concluding chapters, we will focus on a multitude of real-world case studies and problems associated with areas such as computer vision, audio analysis and natural language

processing (NLP). By the end of this book, you will be able to implement both DL and transfer learning principles in your own systems. What you will learn Set up your own DL environment with graphics processing unit (GPU) and Cloud support Delve into transfer learning principles with ML and DL models Explore various DL architectures, including CNN, LSTM, and capsule networks Learn about data and network representation and loss functions Get to grips with models and strategies in transfer learning Walk through potential challenges in building complex transfer learning models from scratch Explore real-world research problems related to computer vision and audio analysis Understand how transfer learning can be leveraged in NLP Who this book is for Hands-On Transfer Learning with Python is for data scientists, machine learning engineers, analysts and developers with an

hands-on-transfer-learning-with-python-implement

interest in data and applying state-of-the-art transfer learning methodologies to solve tough real-world problems. Basic proficiency in machine learning and Python is required.

Hands-On Natural Language Processing with Python - Rajesh Arumugam

2018-07-18

Foster your NLP applications with the help of deep learning, NLTK, and TensorFlow Key Features Weave neural networks into linguistic applications across various platforms Perform NLP tasks and train its models using NLTK and TensorFlow Boost your NLP models with strong deep learning architectures such as CNNs and RNNs Book Description Natural language processing (NLP) has found its application in various domains, such as web search, advertisements, and customer services, and with the help of deep learning, we can enhance its performances in these areas. Hands-On Natural Language Processing with

42/50

Downloaded from info.ucel.edu.ar on by @guest

Python teaches you how to leverage deep learning models for performing various NLP tasks, along with best practices in dealing with today's NLP challenges. To begin with, you will understand the core concepts of NLP and deep learning, such as Convolutional Neural Networks (CNNs), recurrent neural networks (RNNs), semantic embedding, Word2vec, and more. You will learn how to perform each and every task of NLP using neural networks, in which you will train and deploy neural networks in your NLP applications. You will get accustomed to using RNNs and CNNs in various application areas, such as text classification and sequence labeling, which are essential in the application of sentiment analysis, customer service chatbots, and anomaly detection. You will be equipped with practical knowledge in order to implement deep learning in your linguistic applications using Python's popular deep learning library,

TensorFlow. By the end of this book, you will be well versed in building deep learning-backed NLP applications, along with overcoming NLP challenges with best practices developed by domain experts. What you will learn Implement semantic embedding of words to classify and find entities Convert words to vectors by training in order to perform arithmetic operations Train a deep learning model to detect classification of tweets and news Implement a question-answer model with search and RNN models Train models for various text classification datasets using CNN Implement WaveNet a deep generative model for producing a natural-sounding voice Convert voice-to-text and text-to-voice Train a model to convert speech-to-text using DeepSpeech Who this book is for Hands-on Natural Language Processing with Python is for you if you are a developer, machine learning or an NLP engineer who wants to build a deep

learning application that leverages NLP techniques. This comprehensive guide is also useful for deep learning users who want to extend their deep learning skills in building NLP applications. All you need is the basics of machine learning and Python to enjoy the book.

Practical Deep Learning for Cloud, Mobile, and Edge - Anirudh Koul 2019-10-14

Whether you're a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher

hands-on-transfer-learning-with-python-implement

Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley's Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

Hands-On Neural Networks - Leonardo De Marchi 2019-05-30

Design and create neural networks with deep learning and artificial intelligence principles using OpenAI Gym, TensorFlow,

44/50

Downloaded from info.ucel.edu.ar on by @guest

and Keras Key Features Explore neural network architecture and understand how it functions Learn algorithms to solve common problems using back propagation and perceptrons Understand how to apply neural networks to applications with the help of useful illustrations Book Description Neural networks play a very important role in deep learning and artificial intelligence (AI), with applications in a wide variety of domains, right from medical diagnosis, to financial forecasting, and even machine diagnostics. Hands-On Neural Networks is designed to guide you through learning about neural networks in a practical way. The book will get you started by giving you a brief introduction to perceptron networks. You will then gain insights into machine learning and also understand what the future of AI could look like. Next, you will study how embeddings can be used to process textual data and the role of long short-term memory

networks (LSTMs) in helping you solve common natural language processing (NLP) problems. The later chapters will demonstrate how you can implement advanced concepts including transfer learning, generative adversarial networks (GANs), autoencoders, and reinforcement learning. Finally, you can look forward to further content on the latest advancements in the field of neural networks. By the end of this book, you will have the skills you need to build, train, and optimize your own neural network model that can be used to provide predictable solutions. What you will learn Learn how to train a network by using backpropagation Discover how to load and transform images for use in neural networks Study how neural networks can be applied to a varied set of applications Solve common challenges faced in neural network development Understand the transfer learning concept to solve tasks using Keras

and Visual Geometry Group (VGG) network Get up to speed with advanced and complex deep learning concepts like LSTMs and NLP Explore innovative algorithms like GANs and deep reinforcement learning Who this book is for If you are interested in artificial intelligence and deep learning and want to further your skills, then this intermediate-level book is for you. Some knowledge of statistics will help you get the most out of this book.

Advanced Deep Learning with Python - Ivan Vasilev 2019-12-12

Gain expertise in advanced deep learning domains such as neural networks, meta-learning, graph neural networks, and memory augmented neural networks using the Python ecosystem Key Features Get to grips with building faster and more robust deep learning architectures Investigate and train convolutional neural network (CNN) models with GPU-accelerated libraries such

as TensorFlow and PyTorch Apply deep neural networks (DNNs) to computer vision problems, NLP, and GANs Book Description In order to build robust deep learning systems, you'll need to understand everything from how neural networks work to training CNN models. In this book, you'll discover newly developed deep learning models, methodologies used in the domain, and their implementation based on areas of application. You'll start by understanding the building blocks and the math behind neural networks, and then move on to CNNs and their advanced applications in computer vision. You'll also learn to apply the most popular CNN architectures in object detection and image segmentation. Further on, you'll focus on variational autoencoders and GANs. You'll then use neural networks to extract sophisticated vector representations of words, before going on to cover various types of recurrent networks,

such as LSTM and GRU. You'll even explore the attention mechanism to process sequential data without the help of recurrent neural networks (RNNs). Later, you'll use graph neural networks for processing structured data, along with covering meta-learning, which allows you to train neural networks with fewer training samples. Finally, you'll understand how to apply deep learning to autonomous vehicles. By the end of this book, you'll have mastered key deep learning concepts and the different applications of deep learning models in the real world. What you will learn

- Cover advanced and state-of-the-art neural network architectures
- Understand the theory and math behind neural networks
- Train DNNs and apply them to modern deep learning problems
- Use CNNs for object detection and image segmentation
- Implement generative adversarial networks (GANs) and variational autoencoders to

hands-on-transfer-learning-with-python-implement

generate new images Solve natural language processing (NLP) tasks, such as machine translation, using sequence-to-sequence models Understand DL techniques, such as meta-learning and graph neural networks Who this book is for This book is for data scientists, deep learning engineers and researchers, and AI developers who want to further their knowledge of deep learning and build innovative and unique deep learning projects. Anyone looking to get to grips with advanced use cases and methodologies adopted in the deep learning domain using real-world examples will also find this book useful. Basic understanding of deep learning concepts and working knowledge of the Python programming language is assumed.

Hands-On Neural Networks - Leonardo De Marchi 2019-05-30

Design and create neural networks with deep learning and artificial intelligence

47/50

Downloaded from info.ucel.edu.ar on by @guest

principles using OpenAI Gym, TensorFlow, and Keras Key Features Explore neural network architecture and understand how it functions Learn algorithms to solve common problems using back propagation and perceptrons Understand how to apply neural networks to applications with the help of useful illustrations Book Description Neural networks play a very important role in deep learning and artificial intelligence (AI), with applications in a wide variety of domains, right from medical diagnosis, to financial forecasting, and even machine diagnostics. Hands-On Neural Networks is designed to guide you through learning about neural networks in a practical way. The book will get you started by giving you a brief introduction to perceptron networks. You will then gain insights into machine learning and also understand what the future of AI could look like. Next, you will study how embeddings can be used to process textual

data and the role of long short-term memory networks (LSTMs) in helping you solve common natural language processing (NLP) problems. The later chapters will demonstrate how you can implement advanced concepts including transfer learning, generative adversarial networks (GANs), autoencoders, and reinforcement learning. Finally, you can look forward to further content on the latest advancements in the field of neural networks. By the end of this book, you will have the skills you need to build, train, and optimize your own neural network model that can be used to provide predictable solutions. What you will learn Learn how to train a network by using backpropagation Discover how to load and transform images for use in neural networks Study how neural networks can be applied to a varied set of applications Solve common challenges faced in neural network development Understand the transfer

learning concept to solve tasks using Keras and Visual Geometry Group (VGG) network Get up to speed with advanced and complex deep learning concepts like LSTMs and NLP Explore innovative algorithms like GANs and deep reinforcement learning Who this book is for If you are interested in artificial intelligence and deep learning and want to further your skills, then this intermediate-level book is for you. Some knowledge of statistics will help you get the most out of this book.

Hands-On Q-Learning with Python -

Nazia Habib 2019-04-19

Leverage the power of reward-based training for your deep learning models with Python Key Features Understand Q-learning algorithms to train neural networks using Markov Decision Process (MDP) Study practical deep reinforcement learning using Q-Networks Explore state-based unsupervised learning for machine learning

models Book Description Q-learning is a machine learning algorithm used to solve optimization problems in artificial intelligence (AI). It is one of the most popular fields of study among AI researchers. This book starts off by introducing you to reinforcement learning and Q-learning, in addition to helping you get familiar with OpenAI Gym as well as libraries such as Keras and TensorFlow. A few chapters into the book, you will gain insights into model-free Q-learning and use deep Q-networks and double deep Q-networks to solve complex problems. This book will guide you in exploring use cases such as self-driving vehicles and OpenAI Gym's CartPole problem. You will also learn how to tune and optimize Q-networks and their hyperparameters. As you progress, you will understand the reinforcement learning approach to solving real-world problems. You will also explore how to use Q-learning

and related algorithms in real-world applications such as scientific research. Toward the end, you'll gain a sense of what's in store for reinforcement learning. By the end of this book, you will be equipped with the skills you need to solve reinforcement learning problems using Q-learning algorithms with OpenAI Gym, Keras, and TensorFlow. What you will learn

- Explore the fundamentals of reinforcement learning and the state-action-reward process
- Understand Markov decision processes
- Get well versed with libraries such as Keras, and TensorFlow
- Create and

- deploy model-free learning and deep Q-learning agents with TensorFlow, Keras, and OpenAI Gym
- Choose and optimize a Q-Network's learning parameters and fine-tune its performance
- Discover real-world applications and use cases of Q-learning

Who this book is for If you are a machine learning developer, engineer, or professional who wants to delve into the deep learning approach for a complex environment, then this is the book for you.

Proficiency in Python programming and basic understanding of decision-making in reinforcement learning is assumed.