
Tensorflow Machine Learning Cookbook

Practical Deep Learning for Cloud, Mobile, and Edge

Over 60 recipes to design, develop, and deploy self-learning AI models using Python

Keras Deep Learning Cookbook

Deep Learning with TensorFlow

Deep Learning with TensorFlow 2 and Keras

Regression, ConvNets, GANs, RNNs, NLP, and more with TensorFlow 2 and the Keras API, 2nd Edition

Solve Complex Neural Net Problems with TensorFlow, H2O and MXNet

Implement machine learning solutions to overcome various computer vision challenges

Python Deep Learning Cookbook

A Guide to Building Deep Learning Systems

Learning TensorFlow.js

Deep Learning

Build 13 real-world projects with advanced numerical computations using the Python ecosystem

TensorFlow Machine Learning Cookbook

Over 60 recipes to build intelligent machine learning systems with the power of Python, 2nd Edition

Deep Learning for Computer Vision

Machine Learning with TensorFlow

R Deep Learning Cookbook

Create powerful machine learning algorithms with TensorFlow

A Case-Based Approach to Understanding Deep Neural Networks

Machine Learning for Cybersecurity Cookbook

A Visual, Interactive Guide to Artificial Intelligence

Building Machine Learning Pipelines

Concepts, Tools, and Techniques to Build Intelligent Systems

TensorFlow 2.0 Computer Vision Cookbook

Deep Learning Illustrated

Neural networks in TensorFlow.js

PyTorch 1.x Reinforcement Learning Cookbook

TensorFlow Machine Learning Cookbook - Second Edition

Over 50 recipes to help you build, train, and deploy learning agents for real-world applications

TensorFlow 1.x Deep Learning Cookbook

TensorFlow 2 Reinforcement Learning Cookbook

Python Machine Learning Cookbook

Learning TensorFlow

Applied Deep Learning

Machine Learning with Python Cookbook
Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow
Deep Learning Cookbook
TensorFlow Machine Learning Projects
Implement Machine Learning and Deep Learning Models with Python

*Tensorflow
Machine
Learning
Cookbook*

*Downloaded
from
ftp.wtvq.com by
guest*

RICHARDSON ASIA

Practical Deep Learning for Cloud, Mobile, and Edge Packt Publishing Ltd
Skip the theory and get the most out of TensorFlow to build production-ready machine learning models
Key Features
Exploit the features of TensorFlow to build and deploy machine learning models
Train neural networks to tackle real-world problems in Computer Vision and NLP
Handy techniques to write production-ready code for your TensorFlow models
Book Description
TensorFlow is an open source software library for Machine Intelligence. The independent recipes in this book will teach you how to use TensorFlow for complex data computations and allow you to dig deeper and gain more insights into your data than ever before. With the help of this book, you will work with recipes for training models, model evaluation, sentiment analysis,

regression analysis, clustering analysis, artificial neural networks, and more. You will explore RNNs, CNNs, GANs, reinforcement learning, and capsule networks, each using Google's machine learning library, TensorFlow. Through real-world examples, you will get hands-on experience with linear regression techniques with TensorFlow. Once you are familiar and comfortable with the TensorFlow ecosystem, you will be shown how to take it to production. By the end of the book, you will be proficient in the field of machine intelligence using TensorFlow. You will also have good insight into deep learning and be capable of implementing machine learning algorithms in real-world scenarios. What you will learn
Become familiar with the basic features of the TensorFlow library
Get to know Linear Regression techniques with TensorFlow
Learn SVMs with hands-on recipes
Implement neural networks to improve predictive modeling
Apply

NLP and sentiment analysis to your data
Master CNN and RNN through practical recipes
Implement the gradient boosted random forest to predict housing prices
Take TensorFlow into production
Who this book is for
If you are a data scientist or a machine learning engineer with some knowledge of linear algebra, statistics, and machine learning, this book is for you. If you want to skip the theory and build production-ready machine learning models using TensorFlow without reading pages and pages of material, this book is for you. Some background in Python programming is assumed.
Downloading the example code for this book
You can download the example code files for all Packt books you have purchased from your account at <http://www.PacktPub.com>.
If you purchased this book elsewhere, you can visit <http://www.PacktPub.com/> s ...

Over 60 recipes to design, develop, and deploy self-learning AI

models using Python

Packt Publishing Ltd

This book is designed to guide you through TensorFlow and how to use it effectively.

Throughout the book, you will work through recipes and get hands-on experience to perform complex data computations, gain insights into your data, and more.

Keras Deep Learning

Cookbook Packt

Publishing Ltd

This book presents practical solutions to the most common

reinforcement learning problems. The recipes in this book will help you understand the fundamental concepts to develop popular RL algorithms. You will gain practical experience in the RL domain using the modern offerings of the PyTorch 1.x library.

Deep Learning with

TensorFlow Packt

Publishing Ltd

Explore machine learning concepts using the latest numerical computing library — TensorFlow — with the help of this comprehensive cookbook
About This Book Your quick guide to implementing TensorFlow in your day-to-day machine learning activities Learn advanced

techniques that bring more accuracy and speed to machine learning
Upgrade your knowledge to the second generation of machine learning with this guide on TensorFlow
Who This Book Is For This book is ideal for data scientists who are familiar with C++ or Python and perform machine learning activities on a day-to-day basis. Intermediate and advanced machine learning implementers who need a quick guide they can easily navigate will find it useful. What You Will Learn Become familiar with the basics of the TensorFlow machine learning library Get to know Linear Regression techniques with TensorFlow Learn SVMs with hands-on recipes Implement neural networks and improve predictions Apply NLP and sentiment analysis to your data Master CNN and RNN through practical recipes Take TensorFlow into production In Detail TensorFlow is an open source software library for Machine Intelligence. The independent recipes in this book will teach you how to use TensorFlow for complex data computations and will let you dig deeper and gain more insights into your data than ever before.

You'll work through recipes on training models, model evaluation, sentiment analysis, regression analysis, clustering analysis, artificial neural networks, and deep learning – each using Google's machine learning library TensorFlow. This guide starts with the fundamentals of the TensorFlow library which includes variables, matrices, and various data sources. Moving ahead, you will get hands-on experience with Linear Regression techniques with TensorFlow. The next chapters cover important high-level concepts such as neural networks, CNN, RNN, and NLP. Once you are familiar and comfortable with the TensorFlow ecosystem, the last chapter will show you how to take it to production. Style and approach This book takes a recipe-based approach where every topic is explicated with the help of a real-world example.
Deep Learning with TensorFlow 2 and Keras MIT Press
Learn how to solve challenging machine learning problems with TensorFlow, Google's revolutionary new software library for deep learning. If you have some

background in basic linear algebra and calculus, this practical book introduces machine-learning fundamentals by showing you how to design systems capable of detecting objects in images, understanding text, analyzing video, and predicting the properties of potential medicines. TensorFlow for Deep Learning teaches concepts through practical examples and helps you build knowledge of deep learning foundations from the ground up. It's ideal for practicing developers with experience designing software systems, and useful for scientists and other professionals familiar with scripting but not necessarily with designing learning algorithms. Learn TensorFlow fundamentals, including how to perform basic computation Build simple learning systems to understand their mathematical foundations Dive into fully connected deep networks used in thousands of applications Turn prototypes into high-quality models with hyperparameter optimization Process images with convolutional neural networks Handle natural language datasets with recurrent neural

networks Use reinforcement learning to solve games such as tic-tac-toe Train deep networks with hardware including GPUs and tensor processing units *Regression, ConvNets, GANs, RNNs, NLP, and more with TensorFlow 2 and the Keras API, 2nd Edition* "O'Reilly Media, Inc." 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 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 **Solve Complex Neural Net Problems with TensorFlow, H2O and MXNet** Packt Publishing Ltd 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 **Implement machine learning solutions to overcome various computer vision challenges** "O'Reilly Media, Inc." Work with advanced topics in deep learning, such as optimization algorithms, hyper-parameter tuning, dropout, and error analysis as well as strategies to address typical problems encountered when training deep neural networks. You'll begin by studying the activation functions mostly with a single neuron (ReLU, sigmoid, and Swish), seeing how to perform linear and logistic regression using TensorFlow, and choosing the right cost function. The next section talks about more complicated neural network architectures with several layers and neurons and explores the problem of

random initialization of weights. An entire chapter is dedicated to a complete overview of neural network error analysis, giving examples of solving problems originating from variance, bias, overfitting, and datasets coming from different distributions. Applied Deep Learning also discusses how to implement logistic regression completely from scratch without using any Python library except NumPy, to let you appreciate how libraries such as TensorFlow allow quick and efficient experiments. Case studies for each method are included to put into practice all theoretical information. You'll discover tips and tricks for writing optimized Python code (for example vectorizing loops with NumPy). What You Will Learn Implement advanced techniques in the right way in Python and TensorFlow Debug and optimize advanced methods (such as dropout and regularization) Carry out error analysis (to realize if one has a bias problem, a variance problem, a data offset problem, and so on) Set up a machine learning project focused on deep learning on a complex dataset Who This Book Is

For Readers with a medium understanding of machine learning, linear algebra, calculus, and basic Python programming.

Python Deep Learning Cookbook TensorFlow Machine Learning Cookbook Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic

language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1 CLASSIFICATION Machine learning basics Classifying with k-Nearest Neighbors Splitting datasets one feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines

Improving classification with the AdaBoost meta algorithm PART 2
FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3
UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify data Simplifying data with the singular value decomposition Big data and MapReduce
A Guide to Building Deep Learning Systems Manning Publications
Companies are spending billions on machine learning projects, but it's money wasted if the models can't be deployed effectively. In this practical guide, Hannes Hapke and Catherine Nelson walk you through the steps of automating a machine learning pipeline using the TensorFlow ecosystem. You'll learn the techniques and tools that will cut deployment time from days to minutes, so that you can focus on developing new models rather than maintaining legacy

systems. Data scientists, machine learning engineers, and DevOps engineers will discover how to go beyond model development to successfully productize their data science projects, while managers will better understand the role they play in helping to accelerate these projects. Understand the steps to build a machine learning pipeline Build your pipeline using components from TensorFlow Extended Orchestrate your machine learning pipeline with Apache Beam, Apache Airflow, and Kubeflow Pipelines Work with data using TensorFlow Data Validation and TensorFlow Transform Analyze a model in detail using TensorFlow Model Analysis Examine fairness and bias in your model performance Deploy models with TensorFlow Serving or TensorFlow Lite for mobile devices Learn privacy-preserving machine learning techniques
[Learning TensorFlow.js](#) "O'Reilly Media, Inc."
Unlock deeper insights into Machine Learning with this vital guide to cutting-edge predictive analytics About This Book Leverage Python's most powerful open-source libraries for

deep learning, data wrangling, and data visualization Learn effective strategies and best practices to improve and optimize machine learning systems and algorithms Ask - and answer - tough questions of your data with robust statistical models, built for a range of datasets Who This Book Is For If you want to find out how to use Python to start answering critical questions of your data, pick up Python Machine Learning - whether you want to get started from scratch or want to extend your data science knowledge, this is an essential and unmissable resource. What You Will Learn Explore how to use different machine learning models to ask different questions of your data Learn how to build neural networks using Keras and Theano Find out how to write clean and elegant Python code that will optimize the strength of your algorithms Discover how to embed your machine learning model in a web application for increased accessibility Predict continuous target outcomes using regression analysis Uncover hidden patterns and structures in data with clustering Organize

data using effective pre-processing techniques Get to grips with sentiment analysis to delve deeper into textual and social media data In Detail Machine learning and predictive analytics are transforming the way businesses and other organizations operate. Being able to understand trends and patterns in complex data is critical to success, becoming one of the key strategies for unlocking growth in a challenging contemporary marketplace. Python can help you deliver key insights into your data – its unique capabilities as a language let you build sophisticated algorithms and statistical models that can reveal new perspectives and answer key questions that are vital for success. Python Machine Learning gives you access to the world of predictive analytics and demonstrates why Python is one of the world's leading data science languages. If you want to ask better questions of data, or need to improve and extend the capabilities of your machine learning systems, this practical data science book is invaluable. Covering a wide range of powerful Python libraries, including

scikit-learn, Theano, and Keras, and featuring guidance and tips on everything from sentiment analysis to neural networks, you'll soon be able to answer some of the most important questions facing you and your organization. Style and approach Python Machine Learning connects the fundamental theoretical principles behind machine learning to their practical application in a way that focuses you on asking and answering the right questions. It walks you through the key elements of Python and its powerful machine learning libraries, while demonstrating how to get to grips with a range of statistical models.

Deep Learning Packt Publishing Ltd Build machine and deep learning systems with the newly released TensorFlow 2 and Keras for the lab, production, and mobile devices Key Features Introduces and then uses TensorFlow 2 and Keras right from the start Teaches key machine and deep learning techniques Understand the fundamentals of deep learning and machine learning through clear explanations and

extensive code samples Book Description Deep Learning with TensorFlow 2 and Keras, Second Edition teaches neural networks and deep learning techniques alongside TensorFlow (TF) and Keras. You'll learn how to write deep learning applications in the most powerful, popular, and scalable machine learning stack available. TensorFlow is the machine learning library of choice for professional applications, while Keras offers a simple and powerful Python API for accessing TensorFlow. TensorFlow 2 provides full Keras integration, making advanced machine learning easier and more convenient than ever before. This book also introduces neural networks with TensorFlow, runs through the main applications (regression, ConvNets (CNNs), GANs, RNNs, NLP), covers two working example apps, and then dives into TF in production, TF mobile, and using TensorFlow with AutoML. What you will learn Build machine learning and deep learning systems with TensorFlow 2 and the Keras API Use Regression analysis, the most popular approach to machine

learning Understand ConvNets (convolutional neural networks) and how they are essential for deep learning systems such as image classifiers Use GANs (generative adversarial networks) to create new data that fits with existing patterns Discover RNNs (recurrent neural networks) that can process sequences of input intelligently, using one part of a sequence to correctly interpret another Apply deep learning to natural human language and interpret natural language texts to produce an appropriate response Train your models on the cloud and put TF to work in real environments Explore how Google tools can automate simple ML workflows without the need for complex modeling Who this book is for This book is for Python developers and data scientists who want to build machine learning and deep learning systems with TensorFlow. Whether or not you have done machine learning before, this book gives you the theory and practice required to use Keras, TensorFlow 2, and AutoML to build machine learning systems. *Build 13 real-world projects with advanced numerical computations*

using the Python ecosystem Packt Publishing Ltd 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

TensorFlow Machine Learning Cookbook Packt Publishing Ltd

Discover powerful ways to effectively solve real-world machine learning problems using key libraries including scikit-learn, TensorFlow, and PyTorch

Key Features

Learn and implement machine learning algorithms in a variety of real-life scenarios

Cover a range of tasks catering to supervised, unsupervised and reinforcement learning techniques

Find easy-to-follow code solutions for tackling common and not-so-common challenges

Book Description

This eagerly anticipated second edition of the popular Python Machine Learning Cookbook will enable you

to adopt a fresh approach to dealing with real-world machine learning and deep learning tasks. With the help of over 100 recipes, you will learn to build powerful machine learning applications using modern libraries from the Python ecosystem. The book will also guide you on how to implement various machine learning algorithms for classification, clustering, and recommendation engines, using a recipe-based approach. With emphasis on practical solutions, dedicated sections in the book will help you to apply supervised and unsupervised learning techniques to real-world problems. Toward the concluding chapters, you will get to grips with recipes that teach you advanced techniques including reinforcement learning, deep neural networks, and automated machine learning. By the end of this book, you will be equipped with the skills you need to apply machine learning techniques and leverage the full capabilities of the Python ecosystem through real-world examples. What you will learn

Use predictive modeling and apply it to

real-world problems

Explore data visualization techniques to interact with your data

Learn how to build a recommendation engine

Understand how to interact with text data and build models to analyze it

Work with speech data and recognize spoken words using Hidden Markov Models

Get well versed with reinforcement learning, automated ML, and transfer learning

Work with image data and build systems for image recognition and biometric face recognition

Use deep neural networks to build an optical character recognition system

Who this book is for

This book is for data scientists, machine learning developers, deep learning enthusiasts and Python programmers who want to solve real-world challenges using machine-learning techniques and algorithms. If you are facing challenges at work and want ready-to-use code solutions to cover key tasks in machine learning and the deep learning domain, then this book is what you need. Familiarity with Python programming and machine learning concepts will be useful.

Over 60 recipes to build intelligent machine learning systems with the power of Python, 2nd Edition O'Reilly Media

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." -Tim Urban, author of *Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning*

Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. *Deep Learning Illustrated* is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn—with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens—presents straightforward analogies to explain what deep learning is, why it has

become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners. Explore new tools that make deep learning models easier to build, use, and improve. Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative

adversarial networks (GANs), deep reinforcement learning, and more. Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects. Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Deep Learning for Computer Vision

O'Reilly Media, Inc.
TensorFlow Machine Learning Cookbook
Packt Publishing Ltd
[Machine Learning with TensorFlow](#)
Addison-Wesley Professional

Deep learning doesn't have to be intimidating. Until recently, this machine-learning method required years of study, but with frameworks such as Keras and Tensorflow, software engineers without a background in machine learning can quickly enter the field. With the recipes in this cookbook, you'll learn how to solve deep-learning problems for classifying and generating text, images, and music. Each chapter consists of several recipes needed to complete a single project, such as training a music recommending system.

Author Douwe Osinga also provides a chapter with half a dozen techniques to help you if you're stuck. Examples are written in Python with code available on GitHub as a set of Python notebooks. You'll learn how to: Create applications that will serve real users Use word embeddings to calculate text similarity Build a movie recommender system based on Wikipedia links Learn how AIs see the world by visualizing their internal state Build a model to suggest emojis for pieces of text Reuse pretrained networks to build an inverse image search service Compare how GANs, autoencoders and LSTMs generate icons Detect music styles and index song collections

[R Deep Learning Cookbook](#) Packt Publishing Ltd

Get well versed with state-of-the-art techniques to tailor training processes and boost the performance of computer vision models using machine learning and deep learning techniques

Key Features

- Develop, train, and use deep learning algorithms for computer vision tasks using TensorFlow 2.x
- Discover practical recipes to overcome various

challenges faced while building computer vision models

Enable machines to gain a human level understanding to recognize and analyze digital images and videos

Book Description

Computer vision is a scientific field that enables machines to identify and process digital images and videos. This book focuses on independent recipes to help you perform various computer vision tasks using TensorFlow. The book begins by taking you through the basics of deep learning for computer vision, along with covering TensorFlow 2.x's key features, such as the Keras and `tf.data.Dataset` APIs. You'll then learn about the ins and outs of common computer vision tasks, such as image classification, transfer learning, image enhancing and styling, and object detection. The book also covers autoencoders in domains such as inverse image search indexes and image denoising, while offering insights into various architectures used in the recipes, such as convolutional neural networks (CNNs), region-based CNNs (R-CNNs), VGGNet, and You Only

Look Once (YOLO). Moving on, you'll discover tips and tricks to solve any problems faced while building various computer vision applications. Finally, you'll delve into more advanced topics such as Generative Adversarial Networks (GANs), video processing, and AutoML, concluding with a section focused on techniques to help you boost the performance of your networks. By the end of this TensorFlow book, you'll be able to confidently tackle a wide range of computer vision problems using TensorFlow 2.x. What you will learn

- Understand how to detect objects using state-of-the-art models such as YOLOv3
- Use AutoML to predict gender and age from images
- Segment images using different approaches such as FCNs and generative models
- Learn how to improve your network's performance using rank-N accuracy, label smoothing, and test time augmentation
- Enable machines to recognize people's emotions in videos and real-time streams
- Access and reuse advanced TensorFlow Hub models to perform image classification and object detection
- Generate captions for images using

CNNs and RNNs Who this book is for This book is for computer vision developers and engineers, as well as deep learning practitioners looking for go-to solutions to various problems that commonly arise in computer vision. You will discover how to employ modern machine learning (ML) techniques and deep learning architectures to perform a plethora of computer vision tasks. Basic knowledge of Python programming and computer vision is required.

[Create powerful machine learning algorithms with TensorFlow](#) Apress

This cookbook will help you to gain a solid understanding of deep reinforcement learning (RL) algorithms with the help of concise, easy-to-follow implementations from scratch. You'll learn how to implement these algorithms with minimal code and develop AI applications to solve real-world and business problems using RL. *A Case-Based Approach to Understanding Deep Neural Networks* Packt Publishing Ltd Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives

readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better over time. You don't have to be a mathematician to use ML: Tools like

Google's TensorFlow library help with complex calculations so you can focus on getting the answers you need. About the book Machine Learning with TensorFlow, Second Edition is a fully revised guide to building machine learning models using Python and TensorFlow. You'll apply core ML concepts to real-world challenges, such as sentiment analysis, text classification, and image recognition. Hands-on examples illustrate neural network techniques for deep speech processing, facial identification, and auto-encoding with CIFAR-10. What's inside Machine Learning with TensorFlow Choosing the best ML approaches Visualizing algorithms with TensorBoard Sharing results with collaborators Running models in Docker About the reader Requires intermediate Python skills and knowledge of general algebraic concepts like vectors and matrices. Examples use the super-stable 1.15.x branch of TensorFlow and TensorFlow 2.x. About the author Chris Mattmann is the Division Manager of the Artificial Intelligence, Analytics, and Innovation Organization at NASA Jet Propulsion Lab. The first edition of this book was

written by Nishant Shukla
with Kenneth Fricklas.

Table of Contents PART 1
- YOUR MACHINE-
LEARNING RIG 1 A
machine-learning odyssey
2 TensorFlow essentials
PART 2 - CORE LEARNING
ALGORITHMS 3 Linear
regression and beyond 4
Using regression for call-
center volume prediction
5 A gentle introduction to
classification 6 Sentiment

classification: Large
movie-review dataset 7
Automatically clustering
data 8 Inferring user
activity from Android
accelerometer data 9
Hidden Markov models 10
Part-of-speech tagging
and word-sense
disambiguation PART 3 -
THE NEURAL NETWORK
PARADIGM 11 A peek into
autoencoders 12 Applying

autoencoders: The
CIFAR-10 image dataset
13 Reinforcement
learning 14 Convolutional
neural networks 15
Building a real-world CNN:
VGG-Face ad VGG-Face
Lite 16 Recurrent neural
networks 17 LSTMs and
automatic speech
recognition 18 Sequence-
to-sequence models for
chatbots 19 Utility
landscape