

Introduction To Tensorflow Python Programming Tutorials

[TensorFlow Pocket Primer](#)
[Mastering Computer Vision with TensorFlow 2.x](#)
[Machine Learning Pocket Reference](#)
[Deep Learning with Python](#)
[Machine Learning for Economics and Finance in TensorFlow 2](#)
[TensorFlow for Machine Intelligence](#)
[Introduction to Deep Learning](#)
[Learning TensorFlow](#)
[Deep Learning for Coders with fastai and PyTorch](#)
[Machine Learning with TensorFlow](#)
[Learning Deep Learning](#)
[Deep Learning Pipeline](#)
[A Hands-On Introduction to Data Science](#)
[Hands-On Convolutional Neural Networks with TensorFlow](#)
[Deep Learning with Python](#)
[AI and Machine Learning for Coders](#)
[Practical Deep Learning](#)
[TensorFlow for Deep Learning](#)
[Programming with TensorFlow](#)
[Deep Learning with TensorFlow](#)
[Applied Deep Learning](#)
[Deep Learning With Python](#)
[Artificial Intelligence with Python](#)
[Learn TensorFlow 2.0](#)
[TensorFlow For Dummies](#)
[AI Crash Course](#)
[Machine Learning and Deep Learning Using Python and TensorFlow](#)
[Python Data Science Handbook](#)
[TensorFlow in Action](#)
[Reinforcement Learning](#)
[Hands-On Machine Learning with scikit-learn and Scientific Python Toolkits](#)
[Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow](#)
[Bayesian Methods for Hackers](#)
[Pro Deep Learning with TensorFlow](#)
[Machine Learning with TensorFlow, Second Edition](#)
[Deep Learning Classifiers with Memristive Networks](#)
[Ultimate Step by Step Guide to Machine Learning Using Python](#)
[Python Machine Learning](#)
[Introduction to Machine Learning with Python](#)
[Python Machine Learning](#)

Introduction To Tensorflow Python Programming Tutorials

Downloaded from ftp.wtvq.com by guest

DALE KENDRICK

TensorFlow Pocket Primer Simon and Schuster

With detailed notes, tables, and examples, this handy reference will help you navigate the basics of structured machine learning. Author Matt Harrison delivers a valuable guide that you can use for additional support during training and as a convenient resource when you dive into your next machine learning project. Ideal for programmers, data scientists, and AI engineers, this book includes an overview of the machine learning process and walks you through classification with structured data. You'll also learn methods for clustering, predicting a continuous value (regression), and reducing dimensionality, among other topics. This pocket reference includes sections that cover: Classification, using the Titanic dataset Cleaning data and dealing with missing data Exploratory data analysis Common preprocessing steps using sample data Selecting features useful to the model Model selection Metrics and classification evaluation Regression examples using k-nearest neighbor, decision trees, boosting, and more Metrics for regression evaluation Clustering Dimensionality reduction Scikit-learn pipelines

Mastering Computer Vision with TensorFlow 2.x O'Reilly Media

Deploy deep learning solutions in production with ease using TensorFlow. You'll also develop the mathematical understanding and intuition required to invent new deep learning architectures and solutions on your own. Pro Deep Learning with TensorFlow provides practical, hands-on expertise so you can learn deep learning from scratch and deploy meaningful deep learning solutions. This book will allow you to get up to speed quickly using TensorFlow and to optimize different deep learning architectures. All of the practical aspects of deep learning that are relevant in any industry are emphasized in this book. You will be able to use the prototypes demonstrated to build new deep learning applications. The code presented in the book is available in the form of iPython notebooks and scripts which allow you to try out examples and extend them in interesting ways. You will be equipped with the mathematical foundation and scientific knowledge to pursue research in this field and give back to the community. What You'll Learn Understand full stack deep learning using TensorFlow and gain a solid mathematical foundation for deep learning Deploy complex deep learning solutions in production using TensorFlow Carry out research on deep learning and perform experiments using TensorFlow Who This Book Is For Data scientists and machine learning professionals, software developers, graduate students, and open source enthusiasts

Machine Learning Pocket Reference Apress

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 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 Python O'Reilly Media

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.

Machine Learning for Economics and Finance in TensorFlow 2 "O'Reilly Media, Inc."

Integrate scikit-learn with various tools such as NumPy, pandas, imbalanced-learn, and scikit-surprise and use it to solve real-world machine learning problems Key FeaturesDelve into machine learning with this comprehensive guide to scikit-learn and scientific PythonMaster the art of data-driven problem-solving with hands-on examplesFoster your theoretical and practical knowledge of supervised and unsupervised machine learning algorithmsBook Description Machine learning is applied everywhere, from business to research and academia, while scikit-learn is a versatile library that is popular among machine learning practitioners. This book serves as a practical guide for anyone looking to provide hands-on machine learning solutions with scikit-learn and Python toolkits. The book begins with an explanation of machine learning concepts and fundamentals, and strikes a balance between theoretical concepts and their applications. Each chapter covers a different set of algorithms, and shows you how to use them to solve real-life problems. You'll also learn about various key supervised and unsupervised machine learning algorithms using practical examples. Whether it is an instance-based learning algorithm, Bayesian estimation, a deep neural network, a tree-based ensemble, or a recommendation system, you'll gain a thorough understanding of its theory and learn when to apply it. As you advance, you'll learn how to deal with unlabeled data and when to use different clustering and anomaly detection algorithms. By the end of this machine

learning book, you'll have learned how to take a data-driven approach to provide end-to-end machine learning solutions. You'll also have discovered how to formulate the problem at hand, prepare required data, and evaluate and deploy models in production. What you will learn

Understand when to use supervised, unsupervised, or reinforcement learning algorithms

Find out how to collect and prepare your data for machine learning tasks

Tackle imbalanced data and optimize your algorithm for a bias or variance tradeoff

Apply supervised and unsupervised algorithms to overcome various machine learning challenges

Employ best practices for tuning your algorithm's hyper parameters

Discover how to use neural networks for classification and regression

Build, evaluate, and deploy your machine learning solutions to production

Who this book is for

This book is for data scientists, machine learning practitioners, and anyone who wants to learn how machine learning algorithms work and to build different machine learning models using the Python ecosystem. The book will help you take your knowledge of machine learning to the next level by grasping its ins and outs and tailoring it to your needs. Working knowledge of Python and a basic understanding of underlying mathematical and statistical concepts is required.

[TensorFlow for Machine Intelligence](#) Apress

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.

[Introduction to Deep Learning](#) Apress

Understand the principles and practices of machine learning and deep learning

This hands-on guide lays out machine learning and deep learning techniques and technologies in a style that is approachable, using just the basic math required. Written by a pair of experts in the field, Machine Learning and Deep Learning Using Python and TensorFlow contains case studies in several industries, including banking, insurance, e-commerce, retail, and healthcare. The book shows how to utilize machine learning and deep learning functions in today's smart devices and apps. You will get download links for datasets, code, and sample projects referred to in the text. Coverage includes:

Machine learning and deep learning concepts

Python programming and statistics fundamentals

Regression and logistic regression

Decision trees

Model selection and cross-validation

Cluster analysis

Random forests and boosting

Artificial neural networks

TensorFlow and Keras

Deep learning hyperparameters

Convolutional neural networks

Recurrent neural networks and long short-term memory

[Learning TensorFlow](#) Machine Learning Mastery

Master Bayesian Inference through Practical Examples and Computation-Without Advanced Mathematical Analysis

Bayesian methods of inference are deeply natural and extremely powerful. However, most discussions of Bayesian inference rely on intensely complex mathematical analyses and artificial examples, making it inaccessible to anyone without a strong mathematical background. Now, though, Cameron Davidson-Pilon introduces Bayesian inference from a computational perspective, bridging theory to practice—freeing you to get results using computing power. Bayesian Methods for Hackers illuminates Bayesian inference through probabilistic programming with the powerful PyMC language and the closely related Python tools NumPy, SciPy, and Matplotlib. Using this approach, you can reach effective solutions in small increments, without extensive mathematical intervention. Davidson-Pilon begins by introducing the concepts underlying Bayesian inference, comparing it with other techniques and guiding you through building and training your first Bayesian model. Next, he introduces PyMC through a series of detailed examples and intuitive explanations that have been refined after extensive user feedback. You'll learn how to use the Markov Chain Monte Carlo algorithm, choose appropriate sample sizes and priors, work with loss functions, and apply Bayesian inference in domains ranging from finance to marketing. Once you've mastered these techniques, you'll constantly turn to this guide for the working PyMC code you need to jumpstart future projects. Coverage includes

- Learning the Bayesian “state of mind” and its practical implications
- Understanding how computers perform Bayesian inference
- Using the PyMC Python library to program Bayesian analyses
- Building and debugging models with PyMC
- Testing your model's “goodness of fit”
- Opening the “black box” of the Markov Chain Monte Carlo algorithm to see how and why it works
- Leveraging the power of the “Law of Large Numbers”
- Mastering key concepts, such as clustering, convergence, autocorrelation, and thinning
- Using loss functions to measure an estimate's weaknesses based on your goals and desired outcomes
- Selecting appropriate priors and understanding how their influence changes with dataset size
- Overcoming the “exploration versus exploitation” dilemma: deciding when “pretty good” is good enough
- Using Bayesian inference to improve A/B testing
- Solving data science problems when only small amounts of data are available

Cameron Davidson-Pilon has worked in many areas of applied mathematics, from the evolutionary dynamics of genes and diseases to stochastic modeling of financial prices. His contributions to the open source community include lifelines, an implementation of survival analysis in Python. Educated at the University of Waterloo and at the Independent University of Moscow, he currently works with the online commerce leader Shopify.

[Deep Learning for Coders with fastai and PyTorch](#) John Wiley & Sons

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in TensorFlow, an open-source machine learning framework. “I find I learn computer science material best by sitting down and writing programs,” the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to TensorFlow and its interface with Python, the widely used programming language. Familiarity with

linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.

[Machine Learning with TensorFlow](#) "O'Reilly Media, Inc."

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

[Learning Deep Learning](#) Simon and Schuster

Delve into neural networks, implement deep learning algorithms, and explore layers of data abstraction with the help of this comprehensive TensorFlow guide

About This Book

Learn how to implement advanced techniques in deep learning with Google's brainchild, TensorFlow

Explore deep neural networks and layers of data abstraction with the help of this comprehensive guide

Real-world contextualization through some deep learning problems concerning research and application

Who This Book Is For

The book is intended for a general audience of people interested in machine learning and machine intelligence. A rudimentary level of programming in one language is assumed, as is a basic familiarity with computer science techniques and technologies, including a basic awareness of computer hardware and algorithms. Some competence in mathematics is needed to the level of elementary linear algebra and calculus. What You Will Learn

Learn about machine learning landscapes along with the historical development and progress of deep learning

Learn about deep machine intelligence and GPU computing with the latest TensorFlow 1.x

Access public datasets and utilize them using TensorFlow to load, process, and transform data

Use TensorFlow on real-world datasets, including images, text, and more

Learn how to evaluate the performance of your deep learning models

Using deep learning for scalable object detection and mobile computing

Train machines quickly to learn from data by exploring reinforcement learning techniques

Explore active areas of deep learning research and applications

In Detail

Deep learning is the step that comes after machine learning, and has more advanced implementations. Machine learning is not just for academics anymore, but is becoming a mainstream practice through wide adoption, and deep learning has taken the front seat. As a data scientist, if you want to explore data abstraction layers, this book will be your guide. This book shows how this can be exploited in the real world with complex raw data using TensorFlow 1.x. Throughout the book, you'll learn how to implement deep learning algorithms for machine learning systems and integrate them into your product offerings, including search, image recognition, and language processing. Additionally, you'll learn how to analyze and improve the performance of deep learning models. This can be done by comparing algorithms against benchmarks, along with machine intelligence, to learn from the information and determine ideal behaviors within a specific context. After finishing the book, you will be familiar with machine learning techniques, in particular the use of TensorFlow for deep learning, and will be ready to apply your knowledge to research or commercial projects. Style and approach

This step-by-step guide will explore common, and not so common, deep neural networks and show how these can be exploited in the real world with complex raw data. With the help of practical examples, you will learn how to implement different types of neural nets to build smart applications related to text, speech, and image data processing.

[Deep Learning Pipeline](#) Packt Publishing Ltd

Unlock the TensorFlow design secrets behind successful deep learning applications!

Deep Learning StackOverflow contributor Thushan Ganegedara teaches you the new features of TensorFlow 2 in this hands-on guide. In TensorFlow in Action you will learn:

Fundamentals of TensorFlow

Implementing deep learning networks

Picking a high-level Keras API for model building with confidence

Writing comprehensive end-to-end data pipelines

Building models for computer vision and natural language processing

Utilizing pretrained NLP models

Recent algorithms including transformers, attention models, and EIMo

In TensorFlow in Action, you'll dig into the newest version of Google's amazing TensorFlow framework as you learn to create incredible deep learning applications. Author Thushan Ganegedara uses quirky stories, practical examples, and behind-the-scenes explanations to demystify concepts otherwise trapped in dense academic papers. As you dive into modern deep learning techniques like transformer and attention models, you'll benefit from the unique insights of a top StackOverflow contributor for deep learning and NLP. About the technology

Google's TensorFlow framework sits at the heart of modern deep learning. Boasting practical features like multi-GPU support, network data visualization, and easy production pipelines using TensorFlow Extended (TFX), TensorFlow provides the most efficient path to professional AI applications. And the Keras library, fully integrated into TensorFlow 2, makes it a snap to build and train even complex models for vision, language, and more. About the book

TensorFlow in Action teaches you to construct, train, and deploy deep learning models using TensorFlow 2. In this practical tutorial, you'll build reusable skill hands-on as you create production-ready applications such as a French-to-English translator and a neural network that can write fiction. You'll appreciate the in-depth explanations that go from DL basics to advanced applications in NLP, image processing, and MLOps, complete with important details that you'll return to reference over and over. What's

inside Covers TensorFlow 2.9 Recent algorithms including transformers, attention models, and EIMo Build on pretrained models Writing end-to-end data pipelines with TFX About the reader For Python programmers with basic deep learning skills. About the author Thushan Ganegedara is a senior ML engineer at Canva and TensorFlow expert. He holds a PhD in machine learning from the University of Sydney. Table of Contents PART 1 FOUNDATIONS OF TENSORFLOW 2 AND DEEP LEARNING 1 The amazing world of TensorFlow 2 TensorFlow 2 3 Keras and data retrieval in TensorFlow 2 4 Dipping toes in deep learning 5 State-of-the-art in deep learning: Transformers PART 2 LOOK MA, NO HANDS! DEEP NETWORKS IN THE REAL WORLD 6 Teaching machines to see: Image classification with CNNs 7 Teaching machines to see better: Improving CNNs and making them confess 8 Telling things apart: Image segmentation 9 Natural language processing with TensorFlow: Sentiment analysis 10 Natural language processing with TensorFlow: Language modeling PART 3 ADVANCED DEEP NETWORKS FOR COMPLEX PROBLEMS 11 Sequence-to-sequence learning: Part 1 12 Sequence-to-sequence learning: Part 2 13 Transformers 14 TensorBoard: Big brother of TensorFlow 15 TFX: MLOps and deploying models with TensorFlow

[A Hands-On Introduction to Data Science](#) Springer

Practical Deep Learning teaches total beginners how to build the datasets and models needed to train neural networks for your own DL projects. If you've been curious about artificial intelligence and machine learning but didn't know where to start, this is the book you've been waiting for. Focusing on the subfield of machine learning known as deep learning, it explains core concepts and gives you the foundation you need to start building your own models. Rather than simply outlining recipes for using existing toolkits, Practical Deep Learning teaches you the why of deep learning and will inspire you to explore further. All you need is basic familiarity with computer programming and high school math—the book will cover the rest. After an introduction to Python, you'll move through key topics like how to build a good training dataset, work with the scikit-learn and Keras libraries, and evaluate your models' performance. You'll also learn: How to use classic machine learning models like k-Nearest Neighbors, Random Forests, and Support Vector Machines How neural networks work and how they're trained How to use convolutional neural networks How to develop a successful deep learning model from scratch You'll conduct experiments along the way, building to a final case study that incorporates everything you've learned. The perfect introduction to this dynamic, ever-expanding field, Practical Deep Learning will give you the skills and confidence to dive into your own machine learning projects.

[Hands-On Convolutional Neural Networks with TensorFlow](#) "O'Reilly Media, Inc."

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

[Deep Learning with Python](#) Apress

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 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 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.

[AI and Machine Learning for Coders](#) Packt Publishing Ltd

Deep learning is the most interesting and powerful machine learning technique right now. Top deep learning libraries are available on the Python ecosystem like Theano and TensorFlow. Tap into their power in a few lines of code using Keras, the best-of-breed applied deep learning library. In this Ebook, learn exactly how to get started and apply deep learning to your own machine learning projects.

[Practical Deep Learning](#) "O'Reilly Media, Inc."

Build your own pipeline based on modern TensorFlow approaches rather than outdated engineering concepts. This book shows you how to build a deep learning pipeline for real-life TensorFlow projects. You'll learn what a pipeline is and how it works so you can build a full application easily and rapidly. Then troubleshoot and overcome basic TensorFlow obstacles to easily create functional apps and deploy well-trained models. Step-by-step and example-oriented instructions help you understand each step of the deep learning pipeline while you apply the most straightforward and effective tools to demonstrative problems and datasets. You'll also develop a deep learning project by preparing data, choosing the model that fits that data, and debugging your model to get the best fit to data all using TensorFlow techniques. Enhance your skills by accessing some of the most powerful recent trends in data science. If you've ever considered building your own image or text-tagging solution or entering a Kaggle contest, Deep Learning Pipeline is for you! What You'll Learn Develop a deep learning project using dataStudy and apply various models to your data Debug and troubleshoot the proper model suited for your data Who This Book Is For Developers, analysts, and data scientists looking to add to or enhance their existing skills by accessing some of the most powerful recent trends in data science. Prior experience in Python or other TensorFlow related languages and mathematics would be helpful.

[TensorFlow for Deep Learning](#) MIT Press

Apply neural network architectures to build state-of-the-art computer vision applications using the Python programming language Key Features Gain a fundamental understanding of advanced computer vision and neural network models in use today Cover tasks such as low-level vision, image classification, and object detection Develop deep learning models on cloud platforms and optimize them using TensorFlow Lite and the OpenVINO toolkit Book Description Computer vision allows machines to gain human-level understanding to visualize, process, and analyze images and videos. This book focuses on using TensorFlow to help you learn advanced computer vision tasks such as image acquisition, processing, and analysis. You'll start with the key principles of computer vision and deep learning to build a solid foundation, before covering neural network architectures and understanding how they work rather than using them as a black box. Next, you'll explore architectures such as VGG, ResNet, Inception, R-CNN, SSD, YOLO, and MobileNet. As you advance, you'll learn to use visual search methods using transfer learning. You'll also cover advanced computer vision concepts such as semantic segmentation, image inpainting with GAN's, object tracking, video segmentation, and action recognition. Later, the book focuses on how machine learning and deep learning concepts can be used to perform tasks such as edge detection and face recognition. You'll then discover how to develop powerful neural network models on your PC and on various cloud platforms. Finally, you'll learn to perform model optimization methods to deploy models on edge devices for real-time inference. By the end of this book, you'll have a solid understanding of computer vision and be able to confidently develop models to automate tasks. What you will learn Explore methods of feature extraction and image retrieval and visualize different layers of the neural network model Use TensorFlow for various visual search methods for real-world scenarios Build neural networks or adjust parameters to optimize the performance of models Understand TensorFlow DeepLab to perform semantic segmentation on images and DCGAN for image inpainting Evaluate your model and optimize and integrate it into your application to operate at scale Get up to speed with techniques for performing manual and automated image annotation Who this book is for This book is for computer vision professionals, image processing professionals, machine learning engineers and AI developers who have some knowledge of machine learning and deep learning and want to build expert-level computer vision applications. In addition to familiarity with TensorFlow, Python knowledge will be required to get started with this book.

[Programming with TensorFlow](#) "O'Reilly Media, Inc."

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

[Deep Learning with TensorFlow](#) Packt Publishing Ltd

If you're looking to make a career move from programmer to AI specialist, this is the ideal place to start. Based on Laurence Moroney's extremely successful AI courses, this introductory book provides a hands-on, code-first approach to help you build confidence while you learn key topics. You'll understand how to implement the most common scenarios in machine learning, such as computer vision, natural language processing (NLP), and sequence modeling for web, mobile, cloud, and embedded runtimes. Most books on machine learning begin with a daunting amount of advanced math. This guide is built on practical lessons that let you work directly with the code. You'll learn: How to build models with TensorFlow using skills that employers desire The basics of machine learning by working with code samples How to implement computer vision, including feature detection in images How to use NLP to tokenize and sequence words and sentences Methods for embedding models in Android and iOS How to serve models over the web and in the cloud with TensorFlow Serving