
Neural Networks And Deep Learning Pdf

Neural Networks And Deep Learning Pdf

Downloaded from <ftp.paragourmet.com>
by Sasha & French

NEURAL NETWORKS AND DEEP LEARNING PDF DOWNLOAD AND INSTALL PDF

Invite to our collection, where you can effortlessly download Neural Networks And Deep Learning Pdf to improve your knowing and research study experience. Our vast collection of PDF documents can give useful academic sources that accommodate various subjects and interests. We comprehend the significance of accessing details quickly and easily, so we aim to make the procedure of **downloading Neural Networks And Deep Learning Pdf PDF** from our system straightforward and convenient. With just a few clicks, you can open a world of expertise from our collection without any challenges. Join us in exploring our comprehensive collection and start your PDF downloads today!

DISCOVERING OUR COMPREHENSIVE COLLECTION INCLUDING NEURAL NETWORKS AND DEEP LEARNING PDF

At our platform, we take satisfaction in our extensive collection of PDF documents consisting of Neural Networks And Deep Learning Pdf that cater to numerous passions and areas of study. Whether you are aiming to broaden your understanding or carrying out study, we have a vast array of PDFs that are sure to satisfy your requirements.

Our PDF files Neural Networks And Deep Learning Pdf are meticulously curated and picked to supply valuable insights and info to our customers. We have teamed up with professionals in different areas to make sure that our collection continues to be up-to-date and relevant.

From scientific research study papers to academic sources, our PDF files cover a large range of subjects and subjects. With easy access to our collection, you can rapidly browse through and discover the PDF Neural Networks And Deep Learning Pdf that interest you one of the most.

Our system is dedicated to offering you with a smooth and reliable means to enhance your understanding and research experience. We comprehend the significance of having trusted and useful sources at your disposal, and that's why our PDF collection is continually growing and broadening.

So whether you're a student, expert or merely interested, discovering our considerable collection of PDF files Neural Networks And Deep Learning Pdf is sure to offer you with valuable understandings and understanding. Start searching today to discover interesting new research study opportunities!

STRAIGHTFORWARD ACTIONS TO DOWNLOADING NEURAL NETWORKS AND DEEP LEARNING PDF PDF

At our system, our team believe in making the procedure of downloading PDF documents Neural Networks And Deep Learning Pdf quick and easy. Right here's exactly how you can access and download and install PDFs completely free:

Step 1: Check out our considerable collection of PDF data to discover the one you require.

Step 2: Click the download button alongside the PDF Neural

Networks And Deep Learning Pdf you wish to conserve.

Action 3: Await the PDF documents Neural Networks And Deep Learning Pdf to download and install to your device. This ought to only take a few secs.

And that's it! You can currently access Neural Networks And Deep Learning Pdf PDF documents offline any time and share it with others if you want.

Our company believe that discovering and investigating need to be a straightforward and accessible experience for all. That's why we provide our service totally free, making certain that you can access the info you require without any challenges.

RAISE YOUR LEARNING AND STUDY

At our system, our company believe that education and learning needs to be accessible to all. That's why we provide a huge collection of PDF downloads consisting of **Neural Networks And Deep Learning Pdf** that satisfy a wide variety of rate of interests and subjects. Our academic sources are best for students, specialists, and anyone looking to expand their expertise.

With our PDF downloads, you can access useful details on different subjects, consisting of background, scientific research, modern technology, and off program Neural Networks And Deep Learning Pdf. Our sources are perfect for research study functions and can aid you strengthen your understanding of complicated topics.

Our library is frequently growing, and we make every effort to include brand-new and relevant web content on a regular basis. With our easy to use interface, you can easily browse our system and uncover the current academic resources.

By downloading and install Neural Networks And Deep Learning Pdf, you can boost your learning and research ventures and obtain valuable understandings that can benefit you in your individual and expert life.

So, what are you waiting for? Start discovering our collection today and unlock a globe of knowledge within your reaches.

CONCLUSION

At our system, we aim to supply a convenient and totally free solution that allows you to download Neural Networks And Deep Learning Pdf from our large library effortlessly. Our easy to use user interface makes certain that you can access the details you need without any complications or barriers.

Whether you're a student, specialist, or just curious, our PDF downloads supply important educational resources that can improve your knowledge and understanding of numerous topics. By discovering our comprehensive collection, you can broaden your discovering and research ventures and elevate your understanding of the globe around you.

So why wait? Begin downloading and install **Neural Networks And Deep Learning Pdf** and start discovering our collection today and unlock a globe of knowledge at your fingertips. Whether you're looking to broaden your horizons or carry out research study, our straightforward and complimentary solution is right here to sustain you every step of the way.

Neural Networks and Deep Learning

A Textbook Springer

This book covers both classical and modern models in deep learning. The primary focus is on the theory and algorithms of deep learning. The theory and algorithms of neural networks are particularly important for understanding important concepts, so that one can understand the important design concepts of neural architectures in different applications. Why do neural networks work? When do they work better than off-the-shelf machine-learning models? When is depth useful? Why is training neural networks so hard? What are the pitfalls? The book is also rich in discussing different applications in order to give the practitioner a flavor of how neural architectures are designed for different types of problems. Applications associated with many different areas like recommender systems, machine translation, image captioning, image classification, reinforcement-learning based gaming, and text analytics are covered. The chapters of this book span three categories: The basics of neural networks: Many traditional machine learning models can be understood as special cases of neural networks. An emphasis is placed in the first two chapters on understanding the relationship between traditional machine learning and neural networks. Support vector machines, linear/logistic regression, singular value decomposition, matrix factorization, and recommender systems are shown to be special cases of neural networks. These methods are studied together with recent feature engineering methods like word2vec. Fundamentals of neural networks: A detailed discussion of training and regularization is provided in Chapters 3 and 4. Chapters 5 and 6 present radial-basis function (RBF) networks and restricted Boltzmann machines. Advanced topics in neural networks: Chapters 7 and 8 discuss recurrent neural networks and convolutional neural networks. Several advanced topics like deep reinforcement learning, neural Turing machines, Kohonen self-organizing maps, and generative adversarial networks are introduced in Chapters 9 and 10. The book is written for graduate students, researchers, and practitioners. Numerous exercises are available along with a solution manual to aid in classroom teaching. Where possible, an application-centric view is highlighted in order to provide an understanding of the practical uses of each class of techniques.

Deep Learning MIT Press

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. "Written by three experts in the field, Deep Learning is the only comprehensive book on the subject." —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such

applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

Machine Learning with Neural Networks

An Introduction for Scientists and Engineers Cambridge University Press

This modern and self-contained book offers a clear and accessible introduction to the important topic of machine learning with neural networks. In addition to describing the mathematical principles of the topic, and its historical evolution, strong connections are drawn with underlying methods from statistical physics and current applications within science and engineering. Closely based around a well-established undergraduate course, this pedagogical text provides a solid understanding of the key aspects of modern machine learning with artificial neural networks, for students in physics, mathematics, and engineering. Numerous exercises expand and reinforce key concepts within the book and allow students to hone their programming skills. Frequent references to current research develop a detailed perspective on the state-of-the-art in machine learning research.

Introduction to Deep Learning and Neural Networks with PythonTM

A Practical Guide Academic Press

Introduction to Deep Learning and Neural Networks with PythonTM: A Practical Guide is an intensive step-by-step guide for neuroscientists to fully understand, practice, and build neural networks. Providing math and PythonTM code examples to clarify neural network calculations, by book's end readers will fully understand how neural networks work starting from the simplest model $Y=X$ and building from scratch. Details and explanations are provided on how a generic gradient descent algorithm works based on mathematical and PythonTM examples, teaching you how to use the gradient descent algorithm to manually perform all calculations in both the forward and backward passes of training a neural network. Examines the practical side of deep learning and neural networks Provides a problem-based approach to building artificial neural networks using real data Describes PythonTM functions and features for neuroscientists Uses a careful tutorial approach to describe implementation of neural networks in PythonTM Features math and code examples (via companion website) with helpful instructions for easy implementation

The Principles of Deep Learning Theory

An Effective Theory Approach to Understanding Neural Networks Cambridge University Press

This textbook establishes a theoretical framework for understanding deep learning models of practical relevance. With an approach that borrows from theoretical physics, Roberts and Yaida provide clear and pedagogical explanations of how realistic deep neural networks actually work. To make results from the

theoretical forefront accessible, the authors eschew the subject's traditional emphasis on intimidating formality without sacrificing accuracy. Straightforward and approachable, this volume balances detailed first-principle derivations of novel results with insight and intuition for theorists and practitioners alike. This self-contained textbook is ideal for students and researchers interested in artificial intelligence with minimal prerequisites of linear algebra, calculus, and informal probability theory, and it can easily fill a semester-long course on deep learning theory. For the first time, the exciting practical advances in modern artificial intelligence capabilities can be matched with a set of effective principles, providing a timeless blueprint for theoretical research in deep learning.

Advanced Applied Deep Learning

Convolutional Neural Networks and Object Detection *Apress*

Develop and optimize deep learning models with advanced architectures. This book teaches you the intricate details and subtleties of the algorithms that are at the core of convolutional neural networks. In *Advanced Applied Deep Learning*, you will study advanced topics on CNN and object detection using Keras and TensorFlow. Along the way, you will look at the fundamental operations in CNN, such as convolution and pooling, and then look at more advanced architectures such as inception networks, resnets, and many more. While the book discusses theoretical topics, you will discover how to work efficiently with Keras with many tricks and tips, including how to customize logging in Keras with custom callback classes, what is eager execution, and how to use it in your models. Finally, you will study how object detection works, and build a complete implementation of the YOLO (you only look once) algorithm in Keras and TensorFlow. By the end of the book you will have implemented various models in Keras and learned many advanced tricks that will bring your skills to the next level. What You Will Learn See how convolutional neural networks and object detection work Save weights and models on disk Pause training and restart it at a later stage Use hardware acceleration (GPUs) in your code Work with the Dataset TensorFlow abstraction and use pre-trained models and transfer learning Remove and add layers to pre-trained networks to adapt them to your specific project Apply pre-trained models such as Alexnet and VGG16 to new datasets Who This Book Is For Scientists and researchers with intermediate-to-advanced Python and machine learning know-how. Additionally, intermediate knowledge of Keras and TensorFlow is expected.

Deep Learning Neural Networks

Design and Case Studies *World Scientific Publishing Company*

Deep Learning Neural Networks is the fastest growing field in machine learning. It serves as a powerful computational tool for solving prediction, decision, diagnosis, detection and decision problems based on a well-defined computational architecture. It has been successfully applied to a broad field of applications ranging from computer security, speech recognition, image and video recognition to industrial fault detection, medical diagnostics and finance. This comprehensive textbook is the first in the new emerging field. Numerous case studies are succinctly demonstrated in the text. It is intended for use as a one-semester graduate-level university text and as a textbook for research and development establishments in industry, medicine and financial research.

Neural Networks with R

Smart models using CNN, RNN, deep learning, and artificial intelligence principles *Packt Publishing Ltd*

Uncover the power of artificial neural networks by implementing them through R code. About This Book Develop a strong background in neural networks with R, to implement them in your applications Build smart systems using the power of deep learning Real-world case studies to illustrate the power of neural network models Who This Book Is For This book is intended for anyone who has a statistical background with knowledge in R and wants to work with neural networks to get better results from complex data. If you are interested in artificial intelligence and deep learning and you want to level up, then this book is what you need! What You Will Learn Set up R packages for neural networks and deep learning Understand the core concepts of artificial neural networks Understand neurons, perceptrons, bias, weights, and activation functions Implement supervised and unsupervised machine learning in R for neural networks Predict and classify data automatically using neural networks Evaluate and fine-tune the models you build. In Detail Neural networks are one of the most fascinating machine learning models for solving complex computational problems efficiently. Neural networks are used to solve wide range of problems in different areas of AI and machine learning. This book explains the niche aspects of neural networking and provides you with foundation to get started with advanced topics. The book begins with neural network design using the neural net package, then you'll build a solid foundation knowledge of how a neural network learns from data, and the principles behind it. This book covers various types of neural network including recurrent neural networks and convoluted neural networks. You will not only learn how to train neural networks, but will also explore generalization of these networks. Later we will delve into combining different neural network models and work with the real-world use cases. By the end of this book, you will learn to implement neural network models in your applications with the help of practical examples in the book. Style and approach A step-by-step guide filled with real-world practical examples.

Deep Learning with JavaScript

Neural networks in TensorFlow.js *Simon and Schuster*

Summary Deep learning has transformed the fields of computer vision, image processing, and natural language applications. Thanks to TensorFlow.js, now JavaScript developers can build deep learning apps without relying on Python or R. Deep Learning with JavaScript shows developers how they can bring DL technology to the web. Written by the main authors of the TensorFlow library, this new book provides fascinating use cases and in-depth instruction for deep learning apps in JavaScript in your browser or on Node. Foreword by Nikhil Thorat and Daniel Smilkov. About the technology Running deep learning applications in the browser or on Node-based backends opens up exciting possibilities for smart web applications. With the TensorFlow.js library, you build and train deep learning models with JavaScript. Offering uncompromising production-quality scalability, modularity, and responsiveness, TensorFlow.js really shines for its portability. Its models run anywhere JavaScript runs, pushing ML farther up the application stack. About the book In *Deep Learning with JavaScript*, you'll learn to use TensorFlow.js to build deep learning models that run directly in the browser. This fast-paced book, written by Google engineers, is practical, engaging, and easy to follow. Through diverse examples featuring text analysis, speech processing, image recognition, and self-learning game AI, you'll master all the basics of deep learning and explore advanced concepts, like retraining existing

models for transfer learning and image generation. What's inside

- Image and language processing in the browser
- Tuning ML models with client-side data
- Text and image creation with generative deep learning
- Source code samples to test and modify

About the reader For JavaScript programmers interested in deep learning. About the author Shanging Cai, Stanley Bileschi and Eric D. Nielsen are software engineers with experience on the Google Brain team, and were crucial to the development of the high-level API of TensorFlow.js. This book is based in part on the classic, Deep Learning with Python by François Chollet.

TOC:

PART 1 - MOTIVATION AND BASIC CONCEPTS

- 1 • Deep learning and JavaScript
- PART 2 - A GENTLE INTRODUCTION TO TENSORFLOW.JS
- 2 • Getting started: Simple linear regression in TensorFlow.js
- 3 • Adding nonlinearity: Beyond weighted sums
- 4 • Recognizing images and sounds using convnets
- 5 • Transfer learning: Reusing pretrained neural networks
- PART 3 - ADVANCED DEEP LEARNING WITH TENSORFLOW.JS
- 6 • Working with data
- 7 • Visualizing data and models
- 8 • Underfitting, overfitting, and the universal workflow of machine learning
- 9 • Deep learning for sequences and text
- 10 • Generative deep learning
- 11 • Basics of deep reinforcement learning
- PART 4 - SUMMARY AND CLOSING WORDS
- 12 • Testing, optimizing, and deploying models
- 13 • Summary, conclusions, and beyond

MATLAB Deep Learning

With Machine Learning, Neural Networks and Artificial Intelligence *Apress*

Get started with MATLAB for deep learning and AI with this in-depth primer. In this book, you start with machine learning fundamentals, then move on to neural networks, deep learning, and then convolutional neural networks. In a blend of fundamentals and applications, MATLAB Deep Learning employs MATLAB as the underlying programming language and tool for the examples and case studies in this book. With this book, you'll be able to tackle some of today's real world big data, smart bots, and other complex data problems. You'll see how deep learning is a complex and more intelligent aspect of machine learning for modern smart data analysis and usage. What You'll Learn Use MATLAB for deep learning Discover neural networks and multi-layer neural networks Work with convolution and pooling layers Build a MNIST example with these layers Who This Book Is For Those who want to learn deep learning using MATLAB. Some MATLAB experience may be useful.

Neural Networks and Deep Learning

Neural Networks and Deep Learning, Deep Learning, Blockchain Blueprint *Createspace Independent Publishing Platform*

Ready to crank up a deep neural network to get your self-driving car pick up the kids from school? Want to add 'Neural Networks' and 'Deep Learning' to your LinkedIn profile? Well, hold on there... Before you embark on your epic journey into the world of deep learning, there is basic theory to march through first! Check out this exceptional bundle of 3 books... This bundle contains 3 books: Book 1: Neural Networks & Deep Learning: Deep Learning explained to your granny - A visual introduction for beginners who want to make their own Deep Learning Neural Network... What you will gain from this book: * A deep understanding of how Deep Learning works * A basics comprehension on how to build a Deep Neural Network from scratch Who this book is for: * Beginners who want to approach the topic, but are too afraid of complex math to start! * Two main Types of Machine Learning Algorithms * A practical example of Unsupervised Learning *

What are Neural Networks? * McCulloch-Pitts's Neuron * Types of activation function * Types of network architectures * Learning processes * Advantages and disadvantages * Let us give a memory to our Neural Network * The example of book writing Software * Deep learning: the ability of learning to learn * How does Deep Learning work? * Main architectures and algorithms * Main types of DNN * Available Frameworks and libraries * Convolutional Neural Networks * Tunnel Vision * Convolution * The right Architecture for a Neural Network * Test your Neural Network * A general overview of Deep Learning * What are the limits of Deep Learning? * Deep Learning: the basics * Layers, Learning paradigms, Training, Validation * Main architectures and algorithms * Models for Deep Learning * Probabilistic graphic models * Restricted Boltzmann Machines * Deep Belief Networks Book2: Deep Learning: Deep Learning explained to your granny - A guide for Beginners... What's Inside? * A general overview of Deep Learning * What are the limits of Deep Learning? * Deep Learning: the basics * Layers, Learning paradigms, Training, Validation * Main architectures and algorithms * Convolutional Neural Networks * Models for Deep Learning * Probabilistic graphic models * Restricted Boltzmann Machines * Deep Belief Networks * Available Frameworks and libraries * TensorFlow Book 3: Blockchain Blueprint: The ultimate guide to understanding blockchain, cryptocurrencies, smart contracts and the future of money The current emerging innovation of this decade may be the connected world of computing relying on blockchain encryption. The attention given to this technology by global giant players suggests that it will become the operational philosophy of the economic system of the future, ranging across many industries. Blockchain can become the solution we needed for speeding up the economy and transactions in order to keep up with our multi-device connected world. In this book, high tech expert Pat Nakamoto answers your questions concerning the future of Blockchain technology along with addressing different major developments linked to it, like Smart Contracts, Fintech and Ethereum. Hit download. Now!

Efficient Processing of Deep Neural Networks *Morgan & Claypool Publishers*

This book provides a structured treatment of the key principles and techniques for enabling efficient processing of deep neural networks (DNNs). DNNs are currently widely used for many artificial intelligence (AI) applications, including computer vision, speech recognition, and robotics. While DNNs deliver state-of-the-art accuracy on many AI tasks, it comes at the cost of high computational complexity. Therefore, techniques that enable efficient processing of deep neural networks to improve metrics—such as energy-efficiency, throughput, and latency—without sacrificing accuracy or increasing hardware costs are critical to enabling the wide deployment of DNNs in AI systems. The book includes background on DNN processing; a description and taxonomy of hardware architectural approaches for designing DNN accelerators; key metrics for evaluating and comparing different designs; features of the DNN processing that are amenable to hardware/algorithm co-design to improve energy efficiency and throughput; and opportunities for applying new technologies. Readers will find a structured introduction to the field as well as a formalization and organization of key concepts from contemporary works that provides insights that may spark new ideas.

Neural Networks and Deep Learning

Strengthening Deep Neural Networks

Making AI Less Susceptible to Adversarial Trickery *"O'Reilly Media, Inc."*

As deep neural networks (DNNs) become increasingly common in real-world applications, the potential to deliberately "fool" them with data that wouldn't trick a human presents a new attack vector. This practical book examines real-world scenarios where DNNs—the algorithms intrinsic to much of AI—are used daily to process image, audio, and video data. Author Katy Warr considers attack motivations, the risks posed by this adversarial input, and methods for increasing AI robustness to these attacks. If you're a data scientist developing DNN algorithms, a security architect interested in how to make AI systems more resilient to attack, or someone fascinated by the differences between artificial and biological perception, this book is for you. Delve into DNNs and discover how they could be tricked by adversarial input. Investigate methods used to generate adversarial input capable of fooling DNNs. Explore real-world scenarios and model the adversarial threat. Evaluate neural network robustness; learn methods to increase resilience of AI systems to adversarial data. Examine some ways in which AI might become better at mimicking human perception in years to come.

Practical Convolutional Neural Networks

Implement advanced deep learning models using Python *Packt Publishing Ltd*

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 ...

Neural Networks and Statistical Learning *Springer Nature*

This book provides a broad yet detailed introduction to neural networks and machine learning in a statistical framework. A single, comprehensive resource for study and further research, it explores the major popular neural network models and statistical learning approaches with examples and exercises and allows readers to gain a practical working understanding of the content. This updated new edition presents recently published results and includes six new chapters that correspond to the recent advances in computational learning theory, sparse coding, deep learning, big data and cloud computing. Each chapter features state-of-the-art descriptions and significant research findings. The topics covered include: • multilayer perceptron; • the Hopfield network; • associative memory models; • clustering models and algorithms; • the radial basis function network; • recurrent neural networks; • nonnegative matrix factorization; • independent component analysis; • probabilistic and Bayesian networks; and • fuzzy sets and logic. Focusing on the prominent accomplishments and their practical aspects, this book provides academic and technical staff, as well as graduate students and researchers with a solid foundation and comprehensive reference on the fields of neural networks, pattern recognition, signal processing, and machine learning.

Neural Networks and Deep Learning

Neural Networks and Deep Learning, Deep Learning Explained to Your Granny *Createspace Independent Publishing Platform*

Ready to crank up a neural network to get your self-driving car pick up the kids from school? Want to add 'Deep Learning' to your LinkedIn profile? Well, hold on there... Before you embark on your

epic journey into the world of deep learning, there is basic theory to march through first! Take a step-by-step journey through the basics of Neural Networks and Deep Learning, made so simple that...even your granny could understand it! What you will gain from this book: * A deep understanding of how a Neural Network and Deep Learning work * A basics comprehension on how to build a Deep Neural Network from scratch Who this book is for: * Beginners who want to approach the topic, but are too afraid of complex math to start! What's Inside? * A brief introduction to Machine Learning * Two main Types of Machine Learning Algorithms * A practical example of Unsupervised Learning * What are Neural Networks? * McCulloch-Pitts's Neuron * Types of activation function * Types of network architectures * Learning processes * Advantages and disadvantages * Let us give a memory to our Neural Network * The example of book writing Software * Deep learning: the ability of learning to learn * How does Deep Learning work? * Main architectures and algorithms * Main types of DNN * Available Frameworks and libraries * Convolutional Neural Networks * Tunnel Vision * Convolution * The right Architecture for a Neural Network * Test your Neural Network * A general overview of Deep Learning * What are the limits of Deep Learning? * Deep Learning: the basics * Layers, Learning paradigms, Training, Validation * Main architectures and algorithms * Models for Deep Learning * Probabilistic graphic models * Restricted Boltzmann Machines * Deep Belief Networks * Available Frameworks and libraries * TensorFlow Hit download. Now!

The Perceptron

A Theory of Statistical Separability in Cognitive Systems (Project Para)

Learning Deep Learning

Theory and Practice of Neural Networks, Computer Vision, Nlp, and Transformers Using Tensorflow *Addison-Wesley Professional*

NVIDIA's Full-Color Guide to Deep Learning: All Students Need to Get Started and Get Results Learning Deep Learning is a complete guide to DL. Illuminating both the core concepts and the hands-on programming techniques needed to succeed, this book suits seasoned developers, data scientists, analysts, but also those with no prior machine learning or statistic experience. After introducing the essential building blocks of deep neural networks, such as artificial neurons and fully connected, convolutional, and recurrent layers, Magnus Ekman shows how to use them to build advanced architectures, including the Transformer. He describes how these concepts are used to build modern networks for computer vision and natural language processing (NLP), including Mask R-CNN, GPT, and BERT. And he explains how a natural language translator and a system generating natural language descriptions of images. Throughout, Ekman provides concise, well-annotated code examples using TensorFlow with Keras. Corresponding PyTorch examples are provided online, and the book thereby covers the two dominating Python libraries for DL used in industry and academia. He concludes with an introduction to neural architecture search (NAS), exploring important ethical issues and providing resources for further learning. Explore and master core concepts: perceptrons, gradient-based learning, sigmoid neurons, and back propagation. See how DL frameworks make it easier to develop more complicated and useful neural networks. Discover how convolutional neural networks (CNNs) revolutionize image classification and analysis. Apply recurrent neural networks (RNNs)

and long short-term memory (LSTM) to text and other variable-length sequences Master NLP with sequence-to-sequence networks and the Transformer architecture Build applications for natural language translation and image captioning

Artificial Intelligence for Humans, Volume 3

Deep Learning and Neural Networks *Createspace Independent Publishing Platform*

Neural networks have been a mainstay of artificial intelligence since its earliest days. Now, exciting new technologies such as deep learning and convolution are taking neural networks in bold new directions. In this book, we will demonstrate the neural networks in a variety of real-world tasks such as image recognition and data science. We examine current neural network technologies, including ReLU activation, stochastic gradient descent, cross-entropy, regularization, dropout, and visualization.

Interpretable Machine Learning *Lulu.com*

Learning Deep Architectures for AI *Now Publishers Inc*

Theoretical results suggest that in order to learn the kind of complicated functions that can represent high-level abstractions (e.g. in vision, language, and other AI-level tasks), one may need deep architectures. Deep architectures are composed of multiple levels of non-linear operations, such as in neural nets with many hidden layers or in complicated propositional formulae re-using many sub-formulae. Searching the parameter space of deep architectures is a difficult task, but learning algorithms such as those for Deep Belief Networks have recently been proposed to tackle this problem with notable success, beating the state-of-the-art in certain areas. This paper discusses the motivations and principles regarding learning algorithms for deep architectures, in particular those exploiting as building blocks unsupervised learning of single-layer models such as Restricted Boltzmann Machines, used to construct deeper models such as Deep Belief Networks.

Neural Networks for Natural Language Processing *IGI Global*

Information in today's advancing world is rapidly expanding and becoming widely available. This eruption of data has made handling it a daunting and time-consuming task. Natural language processing (NLP) is a method that applies linguistics and algorithms to large amounts of this data to make it more valuable. NLP improves the interaction between humans and computers, yet there remains a lack of research that focuses on the practical implementations of this trending approach. Neural Networks for Natural Language Processing is a collection of innovative research on the methods and applications of linguistic information processing and its computational properties. This publication will support readers with performing sentence classification and language generation using neural networks, apply deep learning models to solve machine translation and conversation problems, and apply deep structured semantic models on information retrieval and natural language applications. While highlighting topics including deep learning, query entity recognition, and information retrieval, this book is ideally designed for research and development professionals, IT specialists, industrialists, technology developers, data analysts, data scientists, academics, researchers, and students seeking current research on the fundamental concepts and techniques of natural language processing.

Deep Learning

A Practitioner's Approach *"O'Reilly Media, Inc."*

Although interest in machine learning has reached a high point, lofty expectations often scuttle projects before they get very far. How can machine learning—especially deep neural networks—make a real difference in your organization? This hands-on guide not only provides the most practical information available on the subject, but also helps you get started building efficient deep learning networks. Authors Adam Gibson and Josh Patterson provide theory on deep learning before introducing their open-source Deeplearning4j (DL4J) library for developing production-class workflows. Through real-world examples, you'll learn methods and strategies for training deep network architectures and running deep learning workflows on Spark and Hadoop with DL4J. Dive into machine learning concepts in general, as well as deep learning in particular Understand how deep networks evolved from neural network fundamentals Explore the major deep network architectures, including Convolutional and Recurrent Learn how to map specific deep networks to the right problem Walk through the fundamentals of tuning general neural networks and specific deep network architectures Use vectorization techniques for different data types with DataVec, DL4J's workflow tool Learn how to use DL4J natively on Spark and Hadoop

Learn Keras for Deep Neural Networks

A Fast-Track Approach to Modern Deep Learning with Python *Apress*

Learn, understand, and implement deep neural networks in a math- and programming-friendly approach using Keras and Python. The book focuses on an end-to-end approach to developing supervised learning algorithms in regression and classification with practical business-centric use-cases implemented in Keras. The overall book comprises three sections with two chapters in each section. The first section prepares you with all the necessary basics to get started in deep learning. Chapter 1 introduces you to the world of deep learning and its difference from machine learning, the choices of frameworks for deep learning, and the Keras ecosystem. You will cover a real-life business problem that can be solved by supervised learning algorithms with deep neural networks. You'll tackle one use case for regression and another for classification leveraging popular Kaggle datasets. Later, you will see an interesting and challenging part of deep learning: hyperparameter tuning; helping you further improve your models when building robust deep learning applications. Finally, you'll further hone your skills in deep learning and cover areas of active development and research in deep learning. At the end of Learn Keras for Deep Neural Networks, you will have a thorough understanding of deep learning principles and have practical hands-on experience in developing enterprise-grade deep learning solutions in Keras. What You'll Learn Master fast-paced practical deep learning concepts with math- and programming-friendly abstractions. Design, develop, train, validate, and deploy deep neural networks using the Keras framework Use best practices for debugging and validating deep learning models Deploy and integrate deep learning as a service into a larger software service or product Extend deep learning principles into other popular frameworks Who This Book Is For Software engineers and data engineers with basic programming skills in any language and who are keen on exploring deep learning for a career move or an enterprise project.

Applied Neural Networks with TensorFlow 2

API Oriented Deep Learning with Python Apress

Implement deep learning applications using TensorFlow while learning the “why” through in-depth conceptual explanations. You’ll start by learning what deep learning offers over other machine learning models. Then familiarize yourself with several technologies used to create deep learning models. While some of these technologies are complementary, such as Pandas, Scikit-Learn, and Numpy—others are competitors, such as PyTorch, Caffe, and Theano. This book clarifies the positions of deep learning and Tensorflow among their peers. You’ll then work on supervised deep learning models to gain applied experience with the technology. A single-layer of multiple perceptrons will be used to build a shallow neural network before turning it into a deep neural network. After showing the structure of the ANNs, a real-life application will be created with Tensorflow 2.0 Keras API. Next, you’ll work on data augmentation and batch normalization methods. Then, the Fashion MNIST dataset will be used to train a CNN. CIFAR10 and Imagenet pre-trained models will be loaded to create already advanced CNNs. Finally, move into theoretical applications and unsupervised learning with auto-encoders and reinforcement learning with tf-agent models. With this book, you’ll delve into applied deep learning practical functions and build a wealth of knowledge about how to use TensorFlow effectively. What You’ll Learn Compare competing technologies and see why TensorFlow is more popular Generate text, image, or sound with GANs Predict the rating or preference a user will give to an item Sequence data with recurrent neural networks Who This Book Is For Data scientists and programmers new to the fields of deep learning and machine learning APIs.

Applied Artificial Intelligence

Neural networks and deep learning with Python and TensorFlow Wolfgang Beer

About This Book Step into the amazing world of Artificial Intelligence and Machine Learning using this compact and easy to understand book. Dive into Neural Networks and Deep Learning and create your own production ready AI models by using TensorFlow and Keras. Work through simple yet insightful examples that will get you up and running with Artificial Intelligence, TensorFlow and Keras in no time. Who This Book Is For This book is for Python developers who want to understand Neural Networks from ground up and build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play around with the code. What You Will Learn The basic structure and functionality of a Neuron The basic math behind the Neural Network learning process See how to build a simple character recognition model from ground up What classification, regression and clustering is How to use TensorFlow to build production ready models Build a first model with the Keras framework How to predict the survival chance for Titanic passengers How to build a simple book recommender How to detect toxic language with an AI model In Detail Artificial Intelligence became one of the hottest topics in the modern economy, where everything is driven by software, network and data. There exists nearly no startup nor traditional business where Artificial Intelligence is not used extensively across many fields such as search engines, image recognition, robotics or finance. This book gives a ground up, step by step introduction about how a Neural Network is used to learn a given function and to make intelligent data-driven decisions. The book explains how to identify typical use-cases such as classification, regression and clustering in terms of practical and

well known use-cases. This book comes with an introduction into the state-of-the-art Google TensorFlow framework that allows developers to roll out their models in production. On top of TensorFlow, the Keras library is used to simplify the design and training of complex deep-learning models. This book comes with multiple examples that show how to apply Artificial Intelligence and Machine Learning models for use-cases such as handwriting recognition, decision making, text analysis and toxic comment identification as well as the use of AI to recommend products to customers.

Deep Learning and Neural Networks: Concepts, Methodologies, Tools, and Applications

Concepts, Methodologies, Tools, and Applications IGI Global

Due to the growing use of web applications and communication devices, the use of data has increased throughout various industries. It is necessary to develop new techniques for managing data in order to ensure adequate usage. Deep learning, a subset of artificial intelligence and machine learning, has been recognized in various real-world applications such as computer vision, image processing, and pattern recognition. The deep learning approach has opened new opportunities that can make such real-life applications and tasks easier and more efficient. Deep Learning and Neural Networks: Concepts, Methodologies, Tools, and Applications is a vital reference source that trends in data analytics and potential technologies that will facilitate insight in various domains of science, industry, business, and consumer applications. It also explores the latest concepts, algorithms, and techniques of deep learning and data mining and analysis. Highlighting a range of topics such as natural language processing, predictive analytics, and deep neural networks, this multi-volume book is ideally designed for computer engineers, software developers, IT professionals, academicians, researchers, and upper-level students seeking current research on the latest trends in the field of deep learning.

Deep Learning and Neural Networks

Concepts, Methodologies, Tools, and Applications Engineering Science Reference

"This book is a vital reference source that trends in data analytics and potential technologies that will facilitate insight in various domains of science, industry, business, and consumer applications. It also explores the latest concepts, algorithms, and techniques of deep learning and data mining and analysis"--

Principles Of Artificial Neural Networks: Basic Designs To Deep Learning (4th Edition) World Scientific

The field of Artificial Neural Networks is the fastest growing field in Information Technology and specifically, in Artificial Intelligence and Machine Learning. This must-have compendium presents the theory and case studies of artificial neural networks. The volume, with 4 new chapters, updates the earlier edition by highlighting recent developments in Deep-Learning Neural Networks, which are the recent leading approaches to neural networks. Uniquely, the book also includes case studies of applications of neural networks — demonstrating how such case studies are designed, executed and how their results are obtained. The title is written for a one-semester graduate or senior-level undergraduate course on artificial neural networks. It is also intended to be a self-study and a reference text for scientists, engineers and for researchers in medicine, finance and data mining.

Deep Learning and Convolutional Neural Networks for Medical Image Computing

Precision Medicine, High Performance and Large-Scale Datasets *Springer*

This book presents a detailed review of the state of the art in deep learning approaches for semantic object detection and segmentation in medical image computing, and large-scale radiology database mining. A particular focus is placed on the application of convolutional neural networks, with the theory supported by practical examples. Features: highlights how the use of deep neural networks can address new questions and protocols, as well as improve upon existing challenges in medical image computing; discusses the insightful research experience of Dr. Ronald M. Summers; presents a comprehensive review of the latest research and literature; describes a range of different methods that make use of deep learning for object or landmark detection tasks in 2D and 3D medical imaging; examines a varied selection of techniques for semantic segmentation using deep learning principles in medical imaging; introduces a novel approach to interleaved text and image deep mining on a large-scale radiology image database.

Deep Learning on Graphs *Cambridge University Press*

A comprehensive text on foundations and techniques of graph neural networks with applications in NLP, data mining, vision and healthcare.

Evolutionary Approach to Machine Learning and Deep Neural Networks

Neuro-Evolution and Gene Regulatory Networks *Springer*

This book provides theoretical and practical knowledge about a methodology for evolutionary algorithm-based search strategy with the integration of several machine learning and deep learning techniques. These include convolutional neural networks, Gröbner bases, relevance vector machines, transfer learning, bagging and boosting methods, clustering techniques (affinity propagation), and belief networks, among others. The development of such tools contributes to better optimizing methodologies. Beginning with the essentials of evolutionary algorithms and covering interdisciplinary research topics, the contents of this book are valuable for different classes of readers: novice, intermediate, and also expert readers from related fields. Following the chapters on introduction and basic methods, Chapter 3 details a new research direction, i.e., neuro-evolution, an evolutionary method for the generation of deep neural networks, and also describes how evolutionary methods are extended in combination with machine learning techniques. Chapter 4 includes novel methods such as particle swarm optimization based on affinity propagation (PSOAP), and transfer learning for differential evolution (TRADE), another machine learning approach for extending differential evolution. The last chapter is dedicated to the state of the art in gene regulatory network (GRN) research as one of the most interesting and active research fields. The author describes an evolving reaction network, which expands the neuro-evolution methodology to produce a type of genetic network suitable for biochemical systems and has succeeded in designing genetic circuits in synthetic biology. The author also presents real-world GRN application to several artificial intelligent tasks, proposing a framework of motion generation by GRNs (MONGERN), which evolves GRNs to operate a real humanoid robot.

Introduction to Deep Learning *MIT Press*

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.

Grokking Deep Learning *Simon and Schuster*

Summary Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning, a branch of artificial intelligence, teaches computers to learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at DeepMind. Previously, Andrew was a researcher and analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform. Table of Contents Introducing deep learning: why you should learn it Fundamental concepts: how do machines learn? Introduction to neural prediction: forward propagation Introduction to neural learning: gradient descent Learning multiple weights at a time: generalizing gradient descent Building your first deep neural network: introduction to backpropagation How to picture neural networks: in your head and on paper Learning signal and ignoring noise: introduction to regularization and batching Modeling probabilities and nonlinearities: activation

functions Neural learning about edges and corners: intro to convolutional neural networks Neural networks that understand language: king - man + woman == ? Neural networks that write like Shakespeare: recurrent layers for variable-length data Introducing automatic optimization: let's build a deep learning framework Learning to write like Shakespeare: long short-term memory Deep learning on unseen data: introducing federated learning Where to go from here: a brief guide

Artificial Intelligence with Python *Packt Publishing Ltd*

Build real-world Artificial Intelligence applications with Python to intelligently interact with the world around you About This Book Step into the amazing world of intelligent apps using this comprehensive guide Enter the world of Artificial Intelligence, explore it, and create your own applications Work through simple yet insightful examples that will get you up and running with Artificial Intelligence in no time Who This Book Is For This book is for Python developers who want to build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play around with the code. It will also be useful for experienced Python programmers who are looking to use Artificial Intelligence techniques in their existing technology stacks. What You Will Learn Realize different classification and regression techniques Understand the concept of clustering and how to use it to automatically segment data See how to build an intelligent recommender system Understand logic programming and how to use it Build automatic speech recognition systems Understand the basics of heuristic search and genetic programming Develop games using Artificial Intelligence Learn how reinforcement learning works Discover how to build intelligent applications centered on images, text, and time series data See how to use deep learning algorithms and build applications based on it In Detail Artificial Intelligence is becoming increasingly relevant in the modern world where everything is driven by technology and data. It is used extensively across many fields such as search engines, image recognition, robotics, finance, and so on. We will explore various real-world scenarios in this book and you'll learn about various algorithms that can be used to build Artificial Intelligence applications. During the course of this book, you will find out how to make informed decisions about what algorithms to use in a given context. Starting from the basics of Artificial Intelligence, you will learn how to develop various building blocks using different data mining techniques. You will see how to implement different algorithms to get the best possible results, and will understand how to apply them to real-world scenarios. If you want to add an intelligence layer to any application that's based on images, text, stock market, or some other form of data, this exciting book on Artificial Intelligence will definitely be your guide! Style and approach This highly practical book will show you how to implement Artificial Intelligence. The book provides multiple examples enabling you to create smart applications to meet the needs of your organization. In every chapter, we explain an algorithm, implement it, and then build a smart application.

Deep Learning Illustrated

A Visual, Interactive Guide to Artificial Intelligence *Addison-Wesley Professional*

"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.

Neural Networks

Deep Learning and Machine Learning Outlined *Independently Published*

Would you achieve more if you could envision your success? A neural network is a computing system made up of a number of simple, highly interconnected processing elements, which process information by their dynamic state response to external inputs. All of this sounds fancy, but what does it mean for computer intelligence, or for the future? In this book, you will find answers to many practical and theoretical questions related to neural networks, from insights about nodes and hidden layers to error spaces, network analyses, and computing influences. Topics will be discussed, such as: What the definition of neural networks encompasses and what all the elements pertaining to them mean. The main advantages of neural networks and how to leverage and apply them. Limitations to neural networks. How neural networks differ from conventional computing systems. Neural Network applications for medical diagnostics, smart computers, artificial intelligence, and forex or stock trading. Troubleshooting tips for when neural networks stop functioning. If you are even in the least interested in computer technology, artificial intelligence, or what the technological future will bring, you need to read this book and get a better understanding of neural networks and their many applications. This book will bring you to the core of how they function and what you can do with them. Add this book to your cart.

An Introduction to Neural Networks *CRC Press*

Though mathematical ideas underpin the study of neural networks, the author presents the fundamentals without the full mathematical apparatus. All aspects of the field are tackled, including artificial neurons as models of their real counterparts; the geometry of network action in pattern space; gradient

descent methods, including back-propagation; associative memory and Hopfield nets; and self-organization and feature maps. The traditionally difficult topic of adaptive resonance theory is clarified within a hierarchical description of its operation. The book also includes several real-world examples to provide a concrete focus. This should enhance its appeal to those involved in the design, construction and management of networks in commercial environments and who wish to improve their understanding of network simulator packages. As a comprehensive and highly accessible introduction to one of the most important topics in cognitive and computer science, this volume should interest a wide range of readers, both students and professionals, in cognitive science, psychology, computer science and electrical engineering.

Deep Learning with PyTorch *Manning Publications*

Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Summary Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Foreword by Soumith Chintala, Cocreator of PyTorch. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Although many deep learning tools use Python, the PyTorch library is truly Pythonic. Instantly familiar to anyone who knows PyData tools like NumPy and scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's excellent for building quick models, and it scales smoothly from laptop to enterprise. Because companies like Apple, Facebook, and JPMorgan Chase rely on PyTorch, it's a great skill to have as you expand your career options. It's easy to get started with PyTorch. It minimizes cognitive overhead without sacrificing the access to advanced features, meaning you can focus on what matters the most -

building and training the latest and greatest deep learning models and contribute to making a dent in the world. PyTorch is also a snap to scale and extend, and it partners well with other Python tooling. PyTorch has been adopted by hundreds of deep learning practitioners and several first-class players like FAIR, OpenAI, FastAI and Purdue. About the book Deep Learning with PyTorch teaches you to create neural networks and deep learning systems with PyTorch. This practical book quickly gets you to work building a real-world example from scratch: a tumor image classifier. Along the way, it covers best practices for the entire DL pipeline, including the PyTorch Tensor API, loading data in Python, monitoring training, and visualizing results. After covering the basics, the book will take you on a journey through larger projects. The centerpiece of the book is a neural network designed for cancer detection. You'll discover ways for training networks with limited inputs and start processing data to get some results. You'll sift through the unreliable initial results and focus on how to diagnose and fix the problems in your neural network. Finally, you'll look at ways to improve your results by training with augmented data, make improvements to the model architecture, and perform other fine tuning. What's inside Training deep neural networks Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Exploring code samples in Jupyter Notebooks About the reader For Python programmers with an interest in machine learning. About the author Eli Stevens had roles from software engineer to CTO, and is currently working on machine learning in the self-driving-car industry. Luca Antiga is cofounder of an AI engineering company and an AI tech startup, as well as a former PyTorch contributor. Thomas Viehmann is a PyTorch core developer and machine learning trainer and consultant. 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