Deep Learning Deep Learning Explained To Your Granny A Guide For Beginners Machine Learning

Thank you for reading **Deep Learning Deep Learning Explained To Your Granny A Guide For Beginners Machine Learning** . Maybe you have knowledge that, people have look hundreds times for their favorite novels like this Deep Learning Deep Learning Explained To Your Granny A Guide For Beginners Machine Learning , but end up in infectious downloads.
Rather than reading a good book with a cup of coffee in the afternoon, instead they are facing with some malicious virus inside their laptop.

Deep Learning Deep Learning Explained To Your Granny A Guide For Beginners Machine Learning is available in our digital library an online access to it is set as public so you can get it instantly.
Our digital library hosts in multiple locations, allowing you to get the most less latency time to download any of our books like this one.
Merely said, the Deep Learning Deep Learning Explained To Your Granny A Guide For Beginners Machine Learning is universally compatible with any devices to read

MATLAB Deep Learning - Phil Kim 2017-06-15
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.

Machine and Deep Learning Algorithms and Applications - Uday Shankar 2022-05-31
This book introduces basic machine learning concepts and applications for a broad audience that includes students, faculty, and industry practitioners. We begin by describing how machine learning provides capabilities to computers and embedded systems to learn from data. A typical machine learning algorithm involves training, and generally the performance of a machine learning model improves with more training data. Deep learning is a sub-area of machine learning that involves extensive use of layers of artificial neural networks typically trained on massive amounts of data. Machine and deep learning methods are often used in contemporary data science tasks to address the growing data sets and detect, cluster, and classify data patterns. Although machine learning commercial interest has grown relatively recently, the roots of machine learning go back to decades ago. We note that nearly all organizations, including industry, government, defense, and health, are using machine learning to address a variety of needs and applications. The machine learning paradigms presented can be broadly divided into the following three categories: supervised learning, unsupervised learning, and semi-supervised learning. Supervised learning algorithms focus on learning a mapping function, and they are trained with supervision on labeled data. Supervised learning is further sub-divided into classification and regression algorithms. Unsupervised learning typically does not have access to ground truth, and often the goal is to learn or uncover the hidden pattern in the data. Through semi-supervised learning, one can effectively utilize a large volume of unlabeled data and a limited amount of labeled data to improve machine learning model performances. Deep learning and neural networks are also covered in this book. Deep neural networks have attracted a lot of interest during the last ten years due to the availability of graphics processing units (GPU) computational power, big data, and new software platforms. They have strong capabilities in terms of learning complex mapping functions for different types of data. We organize the book as follows. The book starts by introducing concepts in supervised, unsupervised, and semi-supervised learning. Several algorithms and their inner workings are presented within these three categories. We then continue with a brief introduction to artificial neural network algorithms and their properties. In addition, we cover an array of applications and provide extensive bibliography. The book ends with a summary of the key machine learning concepts.

TensorFlow for Deep Learning - Bharath Ramsundar 2018-03-01
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 from the Basics - Koki Saitoh 2021-03-08
Discover ways to implement various deep learning algorithms by leveraging Python and other technologies Key FeaturesLearn deep learning models through several activitiesBegin with simple machine learning problems, and finish by building a complex system of your ownTeach your machines to see by mastering the technologies required for image recognitionBook Description Deep learning is rapidly becoming the most preferred way of solving data problems. This is thanks, in part, to its huge variety of mathematical algorithms and their ability to find patterns that are otherwise invisible to us. Deep Learning from the Basics begins with a fast-paced introduction to deep learning with Python, its definition, characteristics, and applications. You'll learn how to use the Python interpreter and the script files in your applications, and utilize NumPy and Matplotlib in your deep learning models. As you progress through the book, you'll discover backpropagation—an efficient way to calculate the gradients of weight parameters—and study multilayer perceptrons and their limitations, before, finally, implementing a three-layer neural network and calculating multidimensional arrays. By the end of the book, you'll have the knowledge to apply the relevant technologies in deep learning. What you will learnUse Python with minimum external sources to implement deep learning programsStudy the various deep learning and neural network theoriesLearn how to determine learning coefficients and the initial values of weightsImplement trends such as Batch Normalization, Dropout, and AdamExplore applications like automatic driving, image generation, and reinforcement learningWho this book is for Deep Learning from the Basics is designed for data scientists, data analysts, and developers who want to use deep learning techniques to develop efficient solutions. This book is ideal for those who want a deeper understanding as well as an overview of the technologies. Some working knowledge of Python is a must. Knowledge of NumPy and pandas will be beneficial, but not essential.

Deep Learning - Josh Patterson 2017-07-28
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

Explainable AI: Interpreting, Explaining and Visualizing Deep Learning - Wojciech Samek 2019-09-10

The development of "intelligent" systems that can take decisions and perform autonomously might lead to faster and more consistent decisions. A limiting factor for a broader adoption of AI technology is the inherent risks that come with giving up human control and oversight to "intelligent" machines. For sensitive tasks involving critical infrastructures and affecting human well-being or health, it is crucial to limit the possibility of improper, non-robust and unsafe decisions and actions. Before deploying an AI system, we see a strong need to validate its behavior, and thus establish guarantees that it will continue to perform as expected when deployed in a real-world environment. In pursuit of that objective, ways for humans to verify the agreement between the AI decision structure and their own ground-truth knowledge have been explored. Explainable AI (XAI) has developed as a subfield of AI, focused on exposing complex AI models to humans in a systematic and interpretable manner. The 22 chapters included in this book provide a timely snapshot of algorithms, theory, and applications of interpretable and explainable AI and AI techniques that have been proposed recently reflecting the current discourse in this field and providing directions of future development. The book is organized in six parts: towards AI transparency; methods for interpreting AI systems; explaining the decisions of AI systems; evaluating interpretability and explanations; applications of explainable AI; and software for explainable AI.

Dive Into Deep Learning - Joanne Quinn 2019-07-15

The leading experts in system change and learning, with their school-based partners around the world, have created this essential companion to their runaway best-seller, Deep Learning: Engage the World Change the World. This hands-on guide provides a roadmap for building capacity in teachers, schools, districts, and systems to design deep learning, measure progress, and assess conditions needed to activate and sustain innovation. Dive Into Deep Learning: Tools for Engagement is rich with resources educators need to construct and drive meaningful deep learning experiences in order to develop the kind of mindset and know-how that is crucial to becoming a problem-solving change agent in our global society. Designed in full color, this easy-to-use guide is loaded with tools, tips, protocols, and real-world examples. It includes: • A framework for deep learning that provides a pathway to develop the six global competencies needed to flourish in a complex world — character, citizenship, collaboration, communication, creativity, and critical thinking. • Learning progressions to help educators analyze student work and measure progress. • Learning design rubrics, templates and examples for incorporating the four elements of learning design: learning partnerships, pedagogical practices, learning environments, and leveraging digital. • Conditions rubrics, teacher self-assessment tools, and planning guides to help educators build, mobilize, and sustain deep learning in schools and districts. Learn about, improve, and expand your world of learning. Put the joy back into learning for students and adults alike. Dive into deep learning to create learning experiences that give purpose, unleash student potential, and transform not only learning, but life itself.

Interpretable Machine Learning - Christoph Molnar 2020

This book is about making machine learning models and their decisions interpretable. After exploring the concepts of interpretability, you will learn about simple, interpretable models such as decision trees, decision rules and linear regression. Later chapters focus on general model-agnostic methods for interpreting black box models like feature importance and accumulated local effects and explaining individual predictions with Shapley values and LIME. All interpretation methods are explained in depth and discussed critically. How do they work under the hood? What are their strengths and weaknesses? How can their outputs be interpreted? This book will enable you to select and correctly apply the interpretation method that is most suitable for your machine learning project.

Deep learning: deep learning explained to your granny – a guide for beginners - PAT NAKAMOTO 2018-07-17

Ready to crank up a deep 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 Deep Learning, made so simple that...even your granny could understand it! 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! 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 Hit download. Now!

Deep Learning with Python - Francois Chollet 2017-11-30

Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance

Deep Learning in Healthcare - Yen-Wei Chen 2019-11-18

This book provides a comprehensive overview of deep learning (DL) in medical and healthcare applications, including the fundamentals and current advances in medical image analysis, state-of-the-art DL methods for

medical image analysis and real-world, deep learning-based clinical computer-aided diagnosis systems. Deep learning (DL) is one of the key techniques of artificial intelligence (AI) and today plays an important role in numerous academic and industrial areas. DL involves using a neural network with many layers (deep structure) between input and output, and its main advantage of is that it can automatically learn data-driven, highly representative and hierarchical features and perform feature extraction and classification on one network. DL can be used to model or simulate an intelligent system or process using annotated training data. Recently, DL has become widely used in medical applications, such as anatomic modelling, tumour detection, disease classification, computer-aided diagnosis and surgical planning. This book is intended for computer science and engineering students and researchers, medical professionals and anyone interested using DL techniques.

Neural Networks and Deep Learning - Charu C. Aggarwal 2018-08-25
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.

Neural Networks and Deep Learning - Pat Nakamoto 2018-01-12
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!
Deep Learning with R - Abhijit Ghatak 2019-04-13
Deep Learning with R introduces deep learning and neural networks using the R programming language. The book builds on the understanding of the theoretical and mathematical constructs and enables the reader to create applications on computer vision, natural language processing and transfer learning. The book starts with an introduction to machine learning and moves on to describe the basic architecture, different activation functions, forward propagation, cross-entropy loss and backward propagation of a simple neural network. It goes on to create different code segments to construct deep neural networks. It discusses in detail the initialization of network parameters, optimization techniques, and some of the common issues surrounding neural networks such as dealing with NaNs and the vanishing/exploding gradient problem. Advanced variants of multilayered perceptrons namely, convolutional neural networks and sequence models are explained, followed by application to different use cases. The book makes extensive use of the Keras and TensorFlow frameworks.
TensorFlow Deep Learning Projects - Alexey Grigorev 2018-03-28
Leverage the power of Tensorflow to design deep learning systems for a variety of real-world scenarios Key Features Build efficient deep learning pipelines using the popular Tensorflow framework Train neural networks such as ConvNets, generative models, and LSTMs Includes projects related to Computer Vision, stock prediction, chatbots and more Book Description TensorFlow is one of the most popular frameworks used for machine learning and, more recently, deep learning. It provides a fast and efficient framework for training different kinds of deep learning models, with very high accuracy. This book is your guide to master deep learning with TensorFlow with the help of 10 real-world projects. TensorFlow Deep Learning Projects starts with setting up the right TensorFlow environment for deep learning. Learn to train different types of deep learning models using TensorFlow, including Convolutional Neural Networks, Recurrent Neural Networks, LSTMs, and Generative Adversarial Networks. While doing so, you will build end-to-end deep learning solutions to tackle different real-world problems in image processing, recommendation systems, stock prediction, and building chatbots, to name a few. You will also develop systems that perform machine translation, and use reinforcement learning techniques to play games. By the end of this book, you will have mastered all the concepts of deep learning and their implementation with TensorFlow, and will be able to build and train your own deep learning models with TensorFlow confidently. What you will learn Set up the TensorFlow environment for deep learning Construct your own ConvNets for effective image processing Use LSTMs for image caption generation Forecast stock prediction accurately with an LSTM architecture Learn what semantic matching is by detecting duplicate Quora questions Set up an AWS instance with TensorFlow to train GANs Train and set up a chatbot to understand and interpret human input Build an AI capable of playing a video game by itself –and win it! Who this book is for This book is for data scientists, machine learning developers as well as deep learning practitioners, who want to build interesting deep learning projects that leverage the power of Tensorflow. Some understanding of machine learning and deep learning, and familiarity with the TensorFlow framework is all you need to get started with this book.
Deep Learning with R - J.J. Allaire 2018-01-22
Summary Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. The book builds your understanding of deep learning through intuitive explanations and practical examples. Continue your journey into the world of deep learning with Deep Learning with R in Motion, a practical, hands-on video course available exclusively at Manning.com (www.manning.com/livevideo/deep-learning-with-r-in-motion). 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. Deep-learning systems now enable previously impossible smart applications, revolutionizing image recognition and natural-language processing, and identifying complex patterns in data. The Keras deep-learning library provides data scientists and developers working in R a state-of-the-art toolset for tackling deep-learning tasks. About the Book Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. Initially written for Python as Deep Learning with Python by Keras creator and Google AI

researcher François Chollet and adapted for R by RStudio founder J. J. Allaire, this book builds your understanding of deep learning through intuitive explanations and practical examples. You'll practice your new skills with R-based applications in computer vision, natural-language processing, and generative models. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image classification and generation Deep learning for text and sequences About the Reader You'll need intermediate R programming skills. No previous experience with machine learning or deep learning is assumed. About the Authors François Chollet is a deep-learning researcher at Google and the author of the Keras library. J.J. Allaire is the founder of RStudio and the author of the R interfaces to TensorFlow and Keras. 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

Deep Learning By Example - Ahmed Menshawy 2018-02-28
Grasp the fundamental concepts of deep learning using Tensorflow in a hands-on manner Key Features Get a first-hand experience of the deep learning concepts and techniques with this easy-to-follow guide Train different types of neural networks using Tensorflow for real-world problems in language processing, computer vision, transfer learning, and more Designed for those who believe in the concept of 'learn by doing', this book is a perfect blend of theory and code examples Book Description Deep learning is a popular subset of machine learning, and it allows you to build complex models that are faster and give more accurate predictions. This book is your companion to take your first steps into the world of deep learning, with hands-on examples to boost your understanding of the topic. This book starts with a quick overview of the essential concepts of data science and machine learning which are required to get started with deep learning. It introduces you to Tensorflow, the most widely used machine learning library for training deep learning models. You will then work on your first deep learning problem by training a deep feed-forward neural network for digit classification, and move on to tackle other real-world problems in computer vision, language processing, sentiment analysis, and more. Advanced deep learning models such as generative adversarial networks and their applications are also covered in this book. By the end of this book, you will have a solid understanding of all the essential concepts in deep learning. With the help of the examples and code provided in this book, you will be equipped to train your own deep learning models with more confidence. What you will learn Understand the fundamentals of deep learning and how it is different from machine learning Get familiarized with Tensorflow, one of the most popular libraries for advanced machine learning Increase the predictive power of your model using feature engineering Understand the basics of deep learning by solving a digit classification problem of MNIST Demonstrate face generation based on the CelebA database, a promising application of generative models Apply deep learning to other domains like language modeling, sentiment analysis, and machine translation Who this book is for This book targets data scientists and machine learning developers who wish to get started with deep learning. If you know what deep learning is but are not quite sure of how to use it, this book will help you as well. An understanding of statistics and data science concepts is required. Some familiarity with Python programming will also be beneficial.

Deep Learning Illustrated - Jon Krohn 2019-08-05
"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 and Deep Learning - Pat Nakamoto 2018-06-30
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!

Deep Learning from the Basics - Koki Saitoh 2021-03-04
Discover ways to implement various deep learning algorithms by leveraging Python and other technologiesKey Features* Learn deep learning models through several activities* Begin with simple machine learning problems, and finish by building a complex system of your own* Teach your machines to see by mastering the technologies required for image recognitionBook DescriptionDeep learning is rapidly becoming the most preferred way of solving data problems. This is thanks, in part, to its huge variety of mathematical algorithms and their ability to find patterns that are otherwise invisible to us.Deep Learning

from the Basics begins with a fast-paced introduction to deep learning with Python, its definition, characteristics, and applications. You'll learn how to use the Python interpreter and the script files in your applications, and utilize NumPy and Matplotlib in your deep learning models. As you progress through the book, you'll discover backpropagation--an efficient way to calculate the gradients of weight parameters--and study multilayer perceptrons and their limitations, before, finally, implementing a three-layer neural network and calculating multidimensional arrays.By the end of the book, you'll have the knowledge to apply the relevant technologies in deep learning.What you will learn* Use Python with minimum external sources to implement deep learning programs* Study the various deep learning and neural network theories* Learn how to determine learning coefficients and the initial values of weights* Implement trends such as Batch Normalization, Dropout, and Adam* Explore applications like automatic driving, image generation, and reinforcement learningWho this book is forDeep Learning from the Basics is designed for data scientists, data analysts, and developers who want to use deep learning techniques to develop efficient solutions. This book is ideal for those who want a deeper understanding as well as an overview of the technologies. Some working knowledge of Python is a must. Knowledge of NumPy and pandas will be beneficial, but not essential.

Inside Deep Learning - Edward Raff 2022-05-31
Journey through the theory and practice of modern deep learning, and apply innovative techniques to solve everyday data problems. In Inside Deep Learning, you will learn how to: Implement deep learning with PyTorch Select the right deep learning components Train and evaluate a deep learning model Fine tune deep learning models to maximize performance Understand deep learning terminology Adapt existing PyTorch code to solve new problems Inside Deep Learning is an accessible guide to implementing deep learning with the PyTorch framework. It demystifies complex deep learning concepts and teaches you to understand the vocabulary of deep learning so you can keep pace in a rapidly evolving field. No detail is skipped--you'll dive into math, theory, and practical applications. Everything is clearly explained in plain English. About the Technology Deep learning doesn't have to be a black box! Knowing how your models and algorithms actually work gives you greater control over your results. And you don't have to be a mathematics expert or a senior data scientist to grasp what's going on inside a deep learning system. This book gives you the practical insight you need to understand and explain your work with confidence. About the Book Inside Deep Learning illuminates the inner workings of deep learning algorithms in a way that even machine learning novices can understand. You'll explore deep learning concepts and tools through plain language explanations, annotated code, and dozens of instantly useful PyTorch examples. Each type of neural network is clearly presented without complex math, and every solution in this book can run using readily available GPU hardware! What's Inside Select the right deep learning components Train and evaluate a deep learning model Fine tune deep learning models to maximize performance Understand deep learning terminology About the Reader For Python programmers with basic machine learning skills. About the Author Edward Raff is a Chief Scientist at Booz Allen Hamilton, and the author of the JSAT machine learning library. Quotes Pick up this book, and you won't be able to put it down. A rich, engaging knowledge base of deep learning math, algorithms, and models--just like the title says! - From the Foreword by Kirk Borne Ph.D., Chief Science Officer, DataPrime.ai The clearest and easiest book for learning deep learning principles and techniques I have ever read. The graphical representations for the algorithms are an eye-opening revelation. - Richard Vaughan, Purple Monkey Collective A great read for anyone interested in understanding the details of deep learning. - Vishwesh Ravi Shrimali, MBRDI.
Deep Learning for Coders with fastai and PyTorch - Jeremy Howard 2020-06-29
Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning

techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

Machine Learning For Dummies - John Paul Mueller 2016-05-31
Your no-nonsense guide to making sense of machine learning Machine learning can be a mind-boggling concept for the masses, but those who are in the trenches of computer programming know just how invaluable it is. Without machine learning, fraud detection, web search results, real-time ads on web pages, credit scoring, automation, and email spam filtering wouldn't be possible, and this is only showcasing just a few of its capabilities. Written by two data science experts, Machine Learning For Dummies offers a much-needed entry point for anyone looking to use machine learning to accomplish practical tasks. Covering the entry-level topics needed to get you familiar with the basic concepts of machine learning, this guide quickly helps you make sense of the programming languages and tools you need to turn machine learning-based tasks into a reality. Whether you're maddened by the math behind machine learning, apprehensive about AI, perplexed by preprocessing data—or anything in between—this guide makes it easier to understand and implement machine learning seamlessly. Grasp how day-to-day activities are powered by machine learning Learn to 'speak' certain languages, such as Python and R, to teach machines to perform pattern-oriented tasks and data analysis Learn to code in R using R Studio Find out how to code in Python using Anaconda Dive into this complete beginner's guide so you are armed with all you need to know about machine learning!

Deep Learning - John D. Kelleher 2019-09-10
An accessible introduction to the artificial intelligence technology that enables computer vision, speech recognition, machine translation, and driverless cars. Deep learning is an artificial intelligence technology that enables computer vision, speech recognition in mobile phones, machine translation, AI games, driverless cars, and other applications. When we use consumer products from Google, Microsoft, Facebook, Apple, or Baidu, we are often interacting with a deep learning system. In this volume in the MIT Press Essential Knowledge series, computer scientist John Kelleher offers an accessible and concise but comprehensive introduction to the fundamental technology at the heart of the artificial intelligence revolution. Kelleher explains that deep learning enables data-driven decisions by identifying and extracting patterns from large datasets; its ability to learn from complex data makes deep learning ideally suited to take advantage of the rapid growth in big data and computational power. Kelleher also explains some of the basic concepts in deep learning, presents a history of advances in the field, and discusses the current state of the art. He describes the most important deep learning architectures, including autoencoders, recurrent neural networks, and long short-term networks, as well as such recent developments as Generative Adversarial Networks and capsule networks. He also provides a comprehensive (and comprehensible) introduction to the two fundamental algorithms in deep learning: gradient descent and backpropagation. Finally, Kelleher considers the future of deep learning—major trends, possible developments, and significant challenges.

Deep Learning Essentials - Anurag Bhardwaj 2018-01-30
Get to grips with the essentials of deep learning by leveraging the power of Python Key Features Your one-stop solution to get started with the essentials of deep learning and neural network modeling Train different kinds of neural networks to tackle various problems in Natural Language Processing, computer vision, speech recognition, and more Covers popular Python libraries such as Tensorflow, Keras, and more, along with tips on training, deploying and optimizing your deep learning models in the best possible manner Book Description Deep Learning a trending topic in the field of Artificial Intelligence today and can be considered to be an advanced form of machine learning, which is quite tricky to master. This book will help you take your first steps in training efficient deep learning models and applying them in various practical scenarios. You will model, train, and deploy different kinds of neural networks such as Convolutional Neural Network, Recurrent Neural Network, and will see some of their applications in real-world domains including computer vision, natural language processing, speech recognition, and so on. You will build practical projects such as chatbots, implement reinforcement learning to build smart games, and

develop expert systems for image captioning and processing. Popular Python library such as TensorFlow is used in this book to build the models. This book also covers solutions for different problems you might come across while training models, such as noisy datasets, small datasets, and more. This book does not assume any prior knowledge of deep learning. By the end of this book, you will have a firm understanding of the basics of deep learning and neural network modeling, along with their practical applications. What you will learn Get to grips with the core concepts of deep learning and neural networks Set up deep learning library such as TensorFlow Fine-tune your deep learning models for NLP and Computer Vision applications Unify different information sources, such as images, text, and speech through deep learning Optimize and fine-tune your deep learning models for better performance Train a deep reinforcement learning model that plays a game better than humans Learn how to make your models get the best out of your GPU or CPU Who this book is for Aspiring data scientists and machine learning experts who have limited or no exposure to deep learning will find this book to be very useful. If you are looking for a resource that gets you up and running with the fundamentals of deep learning and neural networks, this book is for you. As the models in the book are trained using the popular Python-based libraries such as Tensorflow and Keras, it would be useful to have sound programming knowledge of Python.

Applied Deep Learning - Umberto Michelucci 2018-09-07
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.

Learn AI with Python - Gaurav Leekha 2021-10-19
Build AI applications using Python to intelligently interact with the world around you. KEY FEATURES ● Covers the practical aspects of Machine Learning and Deep Learning concepts with the help of this example-rich guide to Python. ● Includes graphical illustrations of Natural Language Processing and its implementation in NLTK. ● Covers deep learning models such as R-CNN and YOLO for object recognition and teaches how to build an image classifier using CNN. DESCRIPTION The book 'Learn AI with Python' is intended to provide you with a thorough understanding of artificial intelligence as well as the tools necessary to create your intelligent applications. This book introduces you to artificial intelligence and walks you through the process of establishing an AI environment on a variety of platforms. It dives into machine learning models and various predictive modeling techniques, including classification, regression, and clustering. Additionally, it provides hands-on experience with logic programming, ASR, neural networks, and natural language processing through real-world examples and fully functional Python implementation. Finally, the book deals with profound models of learning such as R-CNN and YOLO. Object detection in images is also explained in detail using Convolutional Neural Networks (CNNs), which are also explained. By the end of this book, you will have a firm grasp of machine learning and deep learning techniques, as well as a steered methodology for formulating and solving related problems. WHAT YOU WILL LEARN ● Learn to implement various machine learning and deep learning algorithms to achieve smart results. ● Understand how ML algorithms can be applied to real-life applications. ● Explore logic programming and learn how to use it practically to solve real-life problems. ● Learn to develop different types of artificial neural networks with Python. ● Understand reinforcement learning and how to build an environment and agents using Python. ● Work with NLTK and build an automatic speech recognition system. WHO THIS BOOK IS FOR This book is for anyone interested in learning about artificial intelligence and putting it into practice with Python. This book is also valuable for intermediate Machine Learning practitioners as a reference guide. Readers should be familiar with the fundamental understanding of Python programming and machine learning techniques. TABLE OF CONTENTS 1. Introduction to AI and Python 2. Machine Learning and Its Algorithms 3. Classification and Regression Using Supervised Learning 4. Clustering Using Unsupervised Learning 5. Solving Problems with Logic Programming 6. Natural Language Processing with Python 7. Implementing Speech Recognition with Python 8. Implementing Artificial Neural Network (ANN) with Python 9. Implementing Reinforcement Learning with Python 10. Implementing Deep Learning and Convolutional Neural Network

Deep Learning in Practice - Mehdi Ghayoumi 2021-12-01
Deep Learning in Practice helps you learn how to develop and optimize a model for your projects using Deep Learning (DL) methods and architectures. Key features: Demonstrates a quick review on Python, NumPy, and TensorFlow fundamentals. Explains and provides examples of deploying TensorFlow and Keras in several projects. Explains the fundamentals of Artificial Neural Networks (ANNs). Presents several examples and applications of ANNs. Learning the most popular DL algorithms features. Explains and provides examples for the DL algorithms that are presented in this book. Analyzes the DL network's parameter and hyperparameters. Reviews state-of-the-art DL examples. Necessary and main steps for DL modeling. Implements a Virtual Assistant Robot (VAR) using DL methods. Necessary and fundamental information to choose a proper DL algorithm. Gives instructions to learn how to optimize your DL model IN PRACTICE. This book is useful for undergraduate and graduate students, as well as practitioners in industry and academia. It will serve as a useful reference for learning deep learning fundamentals and implementing a deep learning model for any project, step by step.

Deep Learning Applications - M. Arif Wani 2020-02-29
This book presents a compilation of selected papers from the 17th IEEE International Conference on Machine Learning and Applications (IEEE ICMLA 2018), focusing on use of deep learning technology in application like game playing, medical applications, video analytics, regression/classification, object detection/recognition and robotic control in industrial environments. It highlights novel ways of using deep neural networks to solve real-world problems, and also offers insights into deep learning architectures and algorithms, making it an essential reference guide for academic researchers, professionals, software engineers in industry, and innovative product developers.

Deep Learning - Ian Goodfellow 2016-11-10
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.

Introduction to Deep Learning - Sandro Skansi 2018-02-04

This textbook presents a concise, accessible and engaging first introduction to deep learning, offering a wide range of connectionist models which represent the current state-of-the-art. The text explores the most popular algorithms and architectures in a simple and intuitive style, explaining the mathematical derivations in a step-by-step manner. The content coverage includes convolutional networks, LSTMs, Word2vec, RBMs, DBNs, neural Turing machines, memory networks and autoencoders. Numerous examples in working Python code are provided throughout the book, and the code is also supplied separately at an accompanying website. Topics and features: introduces the fundamentals of machine learning, and the mathematical and computational prerequisites for deep learning; discusses feed-forward neural networks, and explores the modifications to these which can be applied to any neural network; examines convolutional neural networks, and the recurrent connections to a feed-forward neural network; describes the notion of distributed representations, the concept of the autoencoder, and the ideas behind language processing with deep learning; presents a brief history of artificial intelligence and neural networks, and reviews interesting open research problems in deep learning and connectionism. This clearly written and lively primer on deep learning is essential reading for graduate and advanced undergraduate students of computer science, cognitive science and mathematics, as well as fields such as linguistics, logic, philosophy, and psychology.

Neural Networks and Deep Learning - Pat Nakamoto 2018-06-30

What's Inside? This includes 3 manuscripts: 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: Big Data: The revolution that is transforming our work, market and world... "Within 2 days we produce the same amount of data generated by at the beginning of the civilization until 2003," said Eric Schmidt in 2010. According to IBM, by 2020 the world will have generated a mass of data on the order of 40 zettabyte (1021Byte). Just think, for example, of digital content such as photos, videos, blogs, posts, and everything that revolves around social networks; only Facebook marks 30 billion pieces of content each month shared by its users. The explosion of social networks, combined with the emergence of smartphones, justifies the fact that one of the recurring terms of recent years in the field of innovation, marketing and IT is "Big Data." The term Big Data indicates data produced in massive quantities, with remarkable rapidity and in the most diverse formats, which require technologies and resources that go far beyond conventional data management and storage systems. In order to obtain

from the use of this data the maximum results in the shortest possible time or even in real time, specific tools with high computing capabilities are necessary. But what does the Big Data phenomenon mean? Is the proliferation of data simply the sign of an increasingly invasive world? Or is there something more to it? Pat Nakamoto will guide you through the discovery of the world of Big data, which, according to experts, in the near future could become the new gold or oil, in what is a real Data Driven economy.

Deep Learning - Andrew Glassner 2021-06-22

A richly-illustrated, full-color introduction to deep learning that offers visual and conceptual explanations instead of equations. You'll learn how to use key deep learning algorithms without the need for complex math. Ever since computers began beating us at chess, they've been getting better at a wide range of human activities, from writing songs and generating news articles to helping doctors provide healthcare. Deep learning is the source of many of these breakthroughs, and its remarkable ability to find patterns hiding in data has made it the fastest growing field in artificial intelligence (AI). Digital assistants on our phones use deep learning to understand and respond intelligently to voice commands; automotive systems use it to safely navigate road hazards; online platforms use it to deliver personalized suggestions for movies and books - the possibilities are endless. Deep Learning: A Visual Approach is for anyone who wants to understand this fascinating field in depth, but without any of the advanced math and programming usually required to grasp its internals. If you want to know how these tools work, and use them yourself, the answers are all within these pages. And, if you're ready to write your own programs, there are also plenty of supplemental Python notebooks in the accompanying Github repository to get you going. The book's conversational style, extensive color illustrations, illuminating analogies, and real-world examples expertly explain the key concepts in deep learning, including: • How text generators create novel stories and articles • How deep learning systems learn to play and win at human games • How image classification systems identify objects or people in a photo • How to think about probabilities in a way that's useful to everyday life • How to use the machine learning techniques that form the core of modern AI Intellectual adventurers of all kinds can use the powerful ideas covered in Deep Learning: A Visual Approach to build intelligent systems that help us better understand the world and everyone who lives in it. It's the future of AI, and this book allows you to fully envision it. Full Color Illustrations

Deep Learning For Dummies - John Paul Mueller 2019-04-15

Take a deep dive into deep learning Deep learning provides the means for discerning patterns in the data that drive online business and social media outlets. Deep Learning for Dummies gives you the information you need to take the mystery out of the topic—and all of the underlying technologies associated with it. In no time, you'll make sense of those increasingly confusing algorithms, and find a simple and safe environment to experiment with deep learning. The book develops a sense of precisely what deep learning can do at a high level and then provides examples of the major deep learning application types. Includes sample code Provides real-world examples within the approachable text Offers hands-on activities to make learning easier Shows you how to use Deep Learning more effectively with the right tools This book is perfect for those who want to better understand the basis of the underlying technologies that we use each and every day.

Grokking Deep Learning - Andrew W. Trask 2019-01-23

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

Principles and Labs for Deep Learning - Shih-Chia Huang 2021-07-06
Principles and Labs for Deep Learning provides the knowledge and techniques needed to help readers design and develop deep learning models. Deep Learning techniques are introduced through theory, comprehensively illustrated, explained through the TensorFlow source code examples, and analyzed through the visualization of results. The structured methods and labs provided by Dr. Huang and Dr. Le enable readers to become proficient in TensorFlow to build deep Convolutional Neural Networks (CNNs) through custom APIs, high-level Keras APIs, Keras Applications, and TensorFlow Hub. Each chapter has one corresponding Lab with step-by-step instruction to help the reader practice and accomplish a specific learning outcome. Deep Learning has been successfully applied in diverse fields such as computer vision, audio processing, robotics, natural language processing, bioinformatics and chemistry. Because of the huge scope of knowledge in Deep Learning, a lot of time is required to understand and deploy useful, working applications, hence the importance of this new resource. Both theory lessons and experiments are included in each chapter to introduce the techniques and provide source code examples to practice using them. All Labs for this book are placed on GitHub to facilitate the download. The book is written based on the assumption that the reader knows basic Python for programming and basic Machine Learning. Introduces readers to the usefulness of neural networks and Deep Learning methods Provides readers with in-depth understanding of the architecture and operation of Deep Convolutional Neural Networks Demonstrates the visualization needed for designing neural networks Provides readers with an in-depth understanding of regression problems, binary classification problems, multi-category classification problems, Variational Auto-Encoder, Generative Adversarial Network, and Object detection

The Deep Learning Revolution - Terrence J. Sejnowski 2018-10-23
How deep learning—from Google Translate to driverless cars to personal cognitive assistants—is changing our lives and transforming every sector of the economy. The deep learning revolution has brought us driverless cars, the greatly improved Google Translate, fluent conversations with Siri and Alexa, and enormous profits from automated trading on the New York Stock Exchange. Deep learning networks can play poker better than professional poker players and defeat a world champion at Go. In this book, Terry Sejnowski explains how deep learning went from being an arcane academic field to a disruptive technology in the information economy. Sejnowski played an important role in the founding of deep learning, as one of a small group of researchers in the 1980s who challenged the prevailing logic-and-symbol based version of AI. The new version of AI Sejnowski and others developed, which became deep learning, is fueled instead by data. Deep networks learn from data in the same way that babies experience the world, starting with fresh eyes and gradually acquiring the skills needed to navigate novel environments. Learning algorithms extract information from raw data; information can be used to create knowledge; knowledge underlies understanding; understanding leads to wisdom. Someday a driverless car will know the road better than

you do and drive with more skill; a deep learning network will diagnose your illness; a personal cognitive assistant will augment your puny human brain. It took nature many millions of years to evolve human intelligence; AI is on a trajectory measured in decades. Sejnowski prepares us for a deep learning future.

Learning Deep Architectures for AI - Yoshua Bengio 2009
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.

Deep Learning with PyTorch - Eli Stevens 2020-08-04
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

Ultimate Step by Step Guide to Deep Learning Using Python - Daneyal Anis 2020-07-19

*Start your Data Science career using Python today!*Are you ready to start your new exciting career? Ready to master artificial intelligence and deep learning concepts?Are you overwhelmed with complexity of the books on this subject?Then let this breezy and fun little book on Python, Machine Learning and Deep Learning models make you a Data Scientist in 7 days!This book continues from where the first book in the series, Ultimate Step by Step Guide to Machine Learning Using Python, left of. In the first book you were introduced to Python concepts such as: -Data Structures like Pandas -Foundational libraries like Numpy, Seaborn and Scikit-Learn-Regression analysis-Classification-Clustering-Association Learning-Dimension ReductionThis book builds on those concepts to expand on Machine Learning algorithms like: -Linear and Logistical regression-Decision tree-Support vector machines (SVM)After that, this book takes you on a journey into Deep Learning and Neural Networks with important concepts and libraries like: -Convolutional and Recurrent Neural Networks-TensorFlow-Keras-PyTorch-Keras-Apache MXNet-Microsoft Cognitive Toolkit (CNTK)The final part of the book covers all foundational concepts that are required for Amazon Web Services (AWS) Certified Machine Learning Specialization by explaining how to deploy your models at scale on Cloud technologies. While AWS is used in the book for illustrative purposes, Microsoft Azure and Google Cloud are also introduced as alternative cloud technologies. After reading this book you will be able to: -Code in Python with confidence-Build new machine learning and deep learning models from scratch-Know how to clean and prepare your data for analytics-Speak confidently about statistical analysis techniquesData Science was ranked the fast-growing field by LinkedIn and Data Scientist is one of the most highly sought after and lucrative careers in the world!If you are on the fence about making the leap to a new and lucrative career, this is the book for you!What sets this book apart from other books on the topic of Python and Machine learning: -Step by step code examples and explanation-Complex concepts explained visually-Real world applicability of the machine learning and deep learning models introducedWhat do I need to get started?You will have a step by step action plan in place once you finish this book and finally feel that you, can master data science and artificial intelligence and start a lucrative and rewarding career! Ready to dive in to the exciting world of Python and Deep Learning?Then scroll up to the top and hit that BUY BUTTON!