

[Book] Deep Learning Deep Learning Explained To Your Granny A Guide For Beginners Machine Learning

Right here, we have countless ebook **deep learning deep learning explained to your granny a guide for beginners machine learning** and collections to check out. We additionally come up with the money for variant types and in addition to type of the books to browse. The enjoyable book, fiction, history, novel, scientific research, as with ease as various new sorts of books are readily genial here.

As this deep learning deep learning explained to your granny a guide for beginners machine learning, it ends going on bodily one of the favored ebook deep learning deep learning explained to your granny a guide for beginners machine learning collections that we have. This is why you remain in the best website to look the unbelievable books to have.

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-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.

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 * 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 Pipeline-Hisham El-Amir 2019-12-20 Build your own pipeline based on modern TensorFlow approaches rather than outdated engineering concepts. This book shows you how to build a deep learning pipeline for real-life TensorFlow projects. You'll learn what a pipeline is and how it works so you can build a full application easily and rapidly. Then troubleshoot and overcome basic TensorFlow obstacles to easily create functional apps and deploy well-trained models. Step-by-step and example-oriented instructions help you understand each step of the deep learning pipeline while you apply the most straightforward and effective tools to demonstrative problems and datasets. You'll also develop a deep learning project by preparing data, choosing the model that fits that data, and debugging your model to get the best fit to data all using Tensorflow techniques. Enhance your skills by accessing some of the most powerful recent trends in data science. If you've ever considered building your own image or text-tagging solution or entering a Kaggle contest, Deep Learning Pipeline is for you! What You'll Learn Develop a deep learning project using data Study and apply various models to your data Debug and troubleshoot the proper model suited for your data Who This Book Is For Developers, analysts, and data scientists looking to add to or enhance their existing skills by accessing some of the most powerful recent trends in data science. Prior experience in Python or other TensorFlow related languages and mathematics would be helpful.

Deep Learning with Python-François Chollet 2017-10-28 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 Fran ois 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-unspeakable speech and image recognition, to near-human accuracy. We went from machines that could 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 Fran ois 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 Fran ois 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

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 2018-06-30 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 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!

Learn Keras for Deep Neural Networks-Jojo Moolayil 2018-12-07 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.

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-deep-learning-explained-to-your-granny-a-guide-for-beginners-machine-learning

Deep Learning and Physics-Akinori Tanaka 2021-04-24 What is deep learning for those who study physics? Is it completely different from physics? Or is it similar? In recent years, machine learning, including deep learning, has begun to be used in various physics studies. Why is that? Is knowing physics useful in machine learning? Conversely, is knowing machine learning useful in physics? This book is devoted to answers of these questions. Starting with basic ideas of physics, neural networks are derived naturally. And you can learn the concepts of deep learning through the words of physics. In fact, the foundation of machine learning can be attributed to physical concepts. Hamiltonians that determine physical systems characterize various machine learning structures. Statistical physics given by Hamiltonians defines machine learning by neural networks. Furthermore, solving inverse problems in physics through machine learning and generalization essentially provides progress and even revolutions in physics. For these reasons, in recent years interdisciplinary research in machine learning and physics has been expanding dramatically. This book is written for anyone who wants to learn, understand, and apply the relationship between deep learning/machine learning and physics. All that is needed to read this book are the basic concepts in physics: energy and Hamiltonians. The concepts of statistical mechanics and the bracket notation of quantum mechanics, which are explained in columns, are used to explain deep learning frameworks. We encourage you to explore this new active field of machine learning and physics, with this book as a map of the continent to be explored.

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

Artificial Intelligence and Deep Learning for Decision Makers-Dr. Jagreet Kaur 2019-12-28 Learn modern-day technologies from modern-day technical giants DESCRIPTION The aim of this book is to help the readers understand the concept of artificial intelligence and deep learning methods and implement them into their businesses and organizations. The first two chapters describe the introduction of the artificial intelligence and deep learning methods. In the first chapter, the concept of human thinking process, starting from the biochemical responses within the structure of neurons to the problem-solving steps through computational thinking skills are discussed. All chapters after the first two should be considered as the study of different technological and Artificial Intelligence giants of current age. These chapters are placed in a way that each chapter could be considered a separate study of a separate company, which includes the achievements of intelligent services currently provided by the company, discussion on the business model of the company towards the use of the deep learning technologies, the advancement of the web services which are incorporated with intelligent capability introduced by company, the efforts of the company in contributing to the development of the artificial intelligence and deep learning research. KEY FEATURES Real-world success and failure stories of artificial intelligence explained Understand concepts of artificial intelligence and deep learning methods Learn how to use artificial intelligence and deep learning methods Know how to prepare dataset and implement models using industry leading Python packages You'll be able to apply and analyze the results produced by the models for prediction WHAT WILL YOU LEARN How to use the algorithms written in the Python programming language to design models and perform predictions in general datasets Understand use cases in different industries related to the implementation of artificial intelligence and deep learning methods Learn the use of potential ideas in artificial intelligence and deep learning methods to improve the operational processes or new products and how services can be produced based on the methods WHO THIS BOOK IS FOR This book is targeted to business and organization leaders, technology enthusiasts, professionals, and managers who seek knowledge of artificial intelligence and deep learning methods. Table of Contents Artificial Intelligence and Deep Learning Data Science for Business Analysis Decision Making Intelligent Computing Strategies By Google Cognitive Learning Services in IBM Watson Advancement web services by Baidu Improved Social Business by Facebook Personalized Intelligent Computing by Apple Cloud Computing Intelligent by Microsoft

Python Machine Learning-Samuel Burns 2019-03-13 You are interested in becoming a machine learning expert but don't know where to start from? Don't worry you don't need a big boring and expensive Textbook. This book is the best guide for you. Get your copy NOW!! Why this guide is the best one for Data Scientist? Here are the reasons:The author has explored everything about machine learning and deep learning right from the basics. A simple language has been used. Many examples have been given, both theoretically and programmatically. Screenshots showing program outputs have been added. The book is written chronologically, in a step-by-step manner. Book Objectives: The Aims and Objectives of the Book: To help you understand the basics of machine learning and deep learning. Understand the various categories of machine learning algorithms. To help you understand how different machine learning algorithms work. You will learn how to implement various machine learning algorithms programmatically in Python. To help you learn how to use Scikit-Learn and TensorFlow Libraries in Python. To help you know how to analyze data programmatically to extract patterns, trends, and relationships between variables. Who this Book is for? Here are the target readers for this book: Anybody who is a complete beginner to machine learning in Python. Anybody who needs to advance their programming skills in Python for machine learning programming and deep learning. Professionals in data science. Professors, lecturers or tutors who are looking to find better ways to explain machine learning to their students in the simplest and easiest way. Students and academicians, especially those focusing on neural networks, machine learning, and deep learning. What do you need for this Book? You are required to have installed the following on your computer: Python 3.X Numpy Pandas Matplotlib The Author guides you on how to install the rest of the Python libraries that are required for machine learning and deep learning. What is inside the book: Getting Started Environment Setup Using Scikit-Learn Linear Regression with Scikit-Learn K-Nearest Neighbors Algorithm K-Means Clustering Support Vector Machines Neural Networks with Scikit-learn Random Forest Algorithm Using TensorFlow Recurrent Neural Networks with TensorFlow Linear Classifier This book will teach you machine learning classifiers using scikit-learn and tensorflow . The book provides a great overview of functions you can use to build a support vector machine, decision tree, perceptron, and k-nearest neighbors. Thanks of this book you will be able to set up a learning pipeline that handles input and output data, pre-processes it, selects meaningful features, and applies a classifier on it. This book offers a lot of insight into machine learning for both beginners, as well as for professionals, who already use some machine learning techniques. Concepts and the background of these concepts are explained clearly in this tutorial.

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 For Dummies-John Paul Mueller 2019-05-14 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.

Explainable AI: Interpreting, Explaining and Visualizing Deep Learning-Wojciech Samek 2019-08-30 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 need 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.

Introduction to Deep Learning-Eugene Charniak 2019-02-19 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.

Interpretable Machine Learning-Christoph Molnar 2019

Deep Learning By Example-Ahmed Menshawy 2018-02-27 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.

AI for People and Business-Alex Castrounis 2019-07-05 If you're an executive, manager, or anyone interested in leveraging AI within your organization, this is your guide. You'll understand exactly what AI is, learn how to identify AI opportunities, and develop and execute a successful AI vision and strategy. Alex Castrounis, business consultant and former IndyCar engineer and race strategist, examines the value of AI and shows you how to develop an AI vision and strategy that benefits both people and business. AI is exciting, powerful, and game changing—but too many AI initiatives end in failure. With this book, you'll explore the risks, considerations, trade-offs, and constraints for pursuing an AI initiative. You'll learn how to create better human experiences and greater business success through winning AI solutions and human-centered products. Use the book's AIPB Framework to conduct end-to-end, goal-driven innovation and value creation with AI Define a goal-aligned AI vision and strategy for stakeholders, including businesses, customers, and users Leverage AI successfully by focusing on concepts such as scientific innovation and AI readiness and maturity Understand the importance of executive leadership for pursuing AI initiatives "A must read for business executives and managers interested in learning about AI and unlocking its benefits. Alex Castrounis has simplified complex topics so that anyone can begin to leverage AI within their organization." - Dan Park, GM & Director, Uber "Alex Castrounis has been at the forefront of helping organizations understand the promise of AI and leverage its benefits, while avoiding the many pitfalls that can derail success. In this essential book, he shares his expertise with the rest of us." - Dean Wampler, Ph.D., VP, Fast Data Engineering at Lightbend

Deep Learning-Josh Patterson 2017-07-28 How can machine learning—especially deep neural networks—make a real difference in your organization? This hands-on guide not only provides practical information, but helps you get started building efficient deep learning networks. The authors provide the fundamentals of deep learning—tuning, parallelization, vectorization, and building pipelines—that are valid for any library before introducing the 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.

Deep Learning with R-Francois Chollet 2018 Introduces deep learning systems using the powerful Keras library and its R language interface. The book builds your understanding of deep learning through intuitive explanations and practical examples.

Artificial Intelligence Engines-James V Stone 2019-04 The brain has always had a fundamental advantage over conventional computers: it can learn. However, a new generation of artificial intelligence algorithms, in the form of deep neural networks, is rapidly eliminating that advantage. Deep neural networks rely on adaptive algorithms to master a wide variety of tasks, including cancer diagnosis, object recognition, speech recognition, robotic control, chess, poker, backgammon and Go, at super-human levels of performance. In this richly illustrated book, key neural network learning algorithms are explained informally first, followed by detailed mathematical analyses. Topics include both historically important neural networks (e.g. perceptrons), and modern deep neural networks (e.g. generative adversarial networks). Online computer programs, collated from open source repositories, give hands-on experience of neural networks, and PowerPoint slides provide support for teaching. Written in an informal style, with a comprehensive glossary, tutorial appendices (e.g. Bayes' theorem), and a list of further readings, this is an ideal introduction to the algorithmic engines of modern artificial intelligence.

Deep Learning with Python-Chao Pan 2016-06-14 ***** BUY NOW (will soon return to 24.77 \$) *****Are you thinking of learning deep Learning using Python? (For Beginners Only) If you are looking for a beginners guide to learn deep learning, in just a few hours, this book is for you. From AI Sciences Publisher Our books may be the best one for beginners; it's a step-by-step guide for any person who wants to start learning Artificial Intelligence and Data Science from scratch. It will help you in preparing a solid foundation and learn any other high-level courses.To get the most out of the concepts that would be covered, readers are advised to adopt a hands on approach, which would lead to better mental representations.Step-by-Step Guide and Visual Illustrations and ExamplesThis book and the accompanying examples, you would be well suited to tackle problems, which pique your interests using machine learning and deep learning models. Book Objectives This book will help you: Have an appreciation for deep learning and an understanding of their fundamental principles. Have an elementary grasp of deep learning concepts and algorithms. Have achieved a technical background in deep learning and neural networks using Python. Target UsersThe book designed for a variety of target audiences. Anyone who is intrigued by how algorithms arrive at predictions but has no previous knowledge of the field. Software developers and engineers with a strong programming background but seeking to break into the field of machine learning. Seasoned professionals in the field of artificial intelligence and deep learning who desire a bird's eye view of current techniques and approaches. What's Inside This Book? Introduction What is Artificial Intelligence, Machine Learning and Deep Learning? Mathematical Foundations of Deep Learning Understanding Machine Learning Models Evaluation of Machine Learning Models: Overfitting, Underfitting, Bias Variance Tradeoff Fully Connected Neural Networks Convolutional Neural Networks Recurrent Neural Networks Generative Adversarial Networks Deep Reinforcement Learning Introduction to Deep Neural Networks with Keras A First Look at Neural Networks in Keras Introduction to Pytorch The Pytorch Deep Learning Framework Your First Neural Network in Pytorch Deep Learning for Computer Vision Build a Convolutional Neural Network Deep Learning for Natural Language Processing Working with Sequential Data Build a Recurrent Neural Network Frequently Asked Questions Q: Is this book for me and do I need programming experience?A: if you want to smash Deep Learning from scratch, this book is for you. Little programming experience is required. If you already wrote a few lines of code and recognize basic programming statements, you'll be OK. Q: Can I have a refund if this book doesn't fit for me?A: Yes, Amazon refund you if you aren't satisfied, for more information about the amazon refund service please go to the amazon help platform. We will also be happy to help you if you send us an email.***** MONEY BACK GUARANTEE BY AMAZON ***** Editorial reviews"This is an excellent book, it is a very good introduction to deep learning and neural networks. The concepts and terminology are clearly explained. The book also points out several good locations on the internet where users can obtain more information. I was extremely happy with this book and I recommend it for all beginners" - Prof. Alain Simon, EDHEC Business School. Statistician and DataScientist.

Machine Intelligence: Demystifying Machine Learning, Neural Networks and Deep Learning-Suresh Samudrala 2019-01-11 Artificial intelligence and machine learning are considered as hot technologies of this century. As these technologies move from research labs to enterprise data centers, the need for skilled professionals is continuously on the rise. This book is intended for IT and business professionals looking to gain proficiency in these technologies but are turned off by the complex mathematical concepts. This book is also useful for students in the area of artificial intelligence and machine learning to gain a conceptual understanding of the algorithms and get an industry perspective. This book is an ideal place to start your journey as -Core equations of machine learning algorithms are explained in plain English using illustrations, data tables and examples -Intuitive meaning of the mathematics behind popular machine learning algorithms explained -Covers classical machine learning, neural networks and deep learning algorithms At a time when the IT industry is focusing on reskilling its vast human resources, Machine intelligence is a very timely publication. It has a simple approach that builds up from basics, which would help software engineers and students looking to learn about the field as well as those who might have started off without the benefit of a structured introduction or sound basics. Highly recommended. - Siddhartha S, Founder and CEO of Intain - Financial technology startup Suresh has written a very accessible book for practitioners. The book has depth yet avoids excessive mathematics. The coverage of the subject is very good and has most of the concepts required for understanding machine learning if someone is looking for depth. For senior management, it will provide a good overview. It is well written. I highly recommend it. - Whee Teck ONG, CEO of Trusted Source and VP of Singapore Computer Society

An Introduction to Machine Learning-Gopinath Rebala 2019-05-07 Just like electricity, Machine Learning will revolutionize our life in many ways - some of which are not even conceivable today. This book provides a thorough conceptual understanding of Machine Learning techniques and algorithms. Many of the mathematical concepts are explained in an intuitive manner. The book starts with an overview of machine learning and the underlying Mathematical and Statistical concepts before moving onto machine learning topics. It gradually builds up the depth, covering many of the present day machine learning algorithms, ending in Deep Learning and Reinforcement Learning algorithms. The book also covers some of the popular Machine Learning applications. The material in this book is agnostic to any specific programming language or hardware so that readers can try these concepts on whichever platforms they are already familiar with. Offers a comprehensive introduction to Machine Learning, while not assuming any prior knowledge of the topic; Provides a complete overview of available techniques and algorithms in conceptual terms, covering various application domains of machine learning; Not tied to any specific software language or hardware implementation.

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-David Feldspar 2018-02-26 Big data can tremendously improve a company's performance. Book 1: All organizations, especially larger ones, face this one dilemma: Sorting and reading the big data. Whether it's a business, NGO, or another organization, patterns and numbers need to be understood to keep management as effective as possible. This book solves all kinds of problems that any entrepreneur, accountant, or manager can use. Book 2: Many company owners or accountants don't look at all the numbers, misread the numbers, or don't know how to scale faster and more effectively. Book 3: The lofty expectations about machine learning and deep studies and projects have skyrocketed, and yet, there is so much left to be said about the methods that trigger the higher-functioning corners of the human neural networks. With so many data and investments on the line, how can we deepen our understanding of these subjects? That is where this guide will take you to the next level. It touches on exactly those problems and methods that optimize your financing and comprehension of the little details that often get overlooked. Book 4: The modern-age machines can deepen and quicken our comprehension of numerous areas and fields. Should we come across such enlightening logic, we better take in all the information we can get, especially if our operating in such a field requires such advanced skills. Book 5: Reinforcement learning has to be one of the most prevalent study techniques in artificial intelligence. It is considered to be a computational approach to learning, in which someone makes attempts to capitalize on the total amount of rewards it obtains when cooperating with a multifaceted, inexact setting. Book 6: A fresh perspective from an educated point of view is what will be presented in this elaborate guide. The readers of this hidden gem will discover how agents and the environment are developed and correspond with each other. Book 7: In a time when everyone is trying to advertise online and offline, the meaning people attach to a word can be significant. With all the unstructured data out there, and the guidance we need to take texts apart and look at them from a logical perspective, analyzing skills have become mandatory to any individual working in marketing, text content creation, or otherwise. Book 8: Become more acquainted with the guiding principles of writing in a certain way, understanding language and tone in specific texts, and using existing models to analyze each and every word in relation to each other.

Machine Learning and Biometrics-Jucheng Yang 2018-08-29 We are entering the era of big data, and machine learning can be used to analyze this deluge of data automatically. Machine learning has been used to solve many interesting and often difficult real-world problems, and the biometrics is one of the leading applications of machine learning. This book introduces some new techniques on biometrics and machine learning, and new proposals of using machine learning techniques for biometrics as well. This book consists of two parts: "Biometrics" and "Machine Learning for Biometrics." Parts I and II contain four and three chapters, respectively. The book is reviewed by editors: Prof. Jucheng Yang, Prof. Dong Sun Park, Prof. Sook Yoon, Dr. Yarui Chen, and Dr. Chuanlei Zhang.

Deep Learning-Li Deng 2014 Provides an overview of general deep learning methodology and its applications to a variety of signal and information processing tasks

Deep Learning with Keras-Antonio Gulli 2017-04-26 Get to grips with the basics of Keras to implement fast and efficient deep-learning models About This Book Implement various deep-learning algorithms in Keras and see how deep-learning can be used in games See how various deep-learning models and practical use-cases can be implemented using Keras A practical, hands-on guide with real-world examples to give you a strong foundation in Keras Who This Book Is For If you are a data scientist with experience in machine learning or an AI programmer with some exposure to neural networks, you will find this book a useful entry point to deep-learning with Keras. A knowledge of Python is required for this book. What You Will Learn Optimize step-by-step functions on a large neural network using the Backpropagation Algorithm Fine-tune a neural network to improve the quality of results Use deep learning for image and audio processing Use Recursive

Neural Tensor Networks (RNTNs) to outperform standard word embedding in special cases Identify problems for which Recurrent Neural Network (RNN) solutions are suitable Explore the process required to implement Autoencoders Evolve a deep neural network using reinforcement learning In Detail This book starts by introducing you to supervised learning algorithms such as simple linear regression, the classical multilayer perceptron and more sophisticated deep convolutional networks. You will also explore image processing with recognition of hand written digit images, classification of images into different categories, and advanced objects recognition with related image annotations. An example of identification of salient points for face detection is also provided. Next you will be introduced to Recurrent Networks, which are optimized for processing sequence data such as text, audio or time series. Following that, you will learn about unsupervised learning algorithms such as Autoencoders and the very popular Generative Adversarial Networks (GAN). You will also explore non-traditional uses of neural networks as Style Transfer. Finally, you will look at Reinforcement Learning and its application to AI game playing, another popular direction of research and application of neural networks. Style and approach This book is an easy-to-follow guide full of examples and real-world applications to help you gain an in-depth understanding of Keras. This book will showcase more than twenty working Deep Neural Networks coded in Python using Keras.

Deep Learning with R-Abhijit Ghatak 2019-06-08 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.

Fundamentals of Deep Learning-Nikhil Buduma 2017-05-25 With the reinvigoration of neural networks in the 2000s, deep learning has become an extremely active area of research, one that's paving the way for modern machine learning. In this practical book, author Nikhil Buduma provides examples and clear explanations to guide you through major concepts of this complicated field. Companies such as Google, Microsoft, and Facebook are actively growing in-house deep-learning teams. For the rest of us, however, deep learning is still a pretty complex and difficult subject to grasp. If you're familiar with Python, and have a background in calculus, along with a basic understanding of machine learning, this book will get you started. Examine the foundations of machine learning and neural networks Learn how to train feed-forward neural networks Use TensorFlow to implement your first neural network Manage problems that arise as you begin to make networks deeper Build neural networks that analyze complex images Perform effective dimensionality reduction using autoencoders Dive deep into sequence analysis to examine language Learn the fundamentals of reinforcement learning

Deep Learning Fundamentals for Beginners-Ai Publishing 2020-01-13 Deep Learning Fundamentals for Beginners Theory behind explained for Beginners Are you looking for a small book to learn fundamentals of Deep Learning fast? Do you need to start learning Deep Learning and Neural Networks from Scratch? This book is for you. This book works as guide to present fundamental concepts, theory, examples related to Deep Learning. This book is written for beginners and novices who want to develop fundamental data science skills and learn how to build models that learn useful information from data. This book will prepare the learner for a career or further learning that involves more advanced topics. It contains an introduction and fundamental concepts used in data science and deep learning. The learner does not need to have any prior knowledge of machine learning or deep learning, but some basic understanding of mathematics is required. This book assumes that you know nothing about deep learning. Its goal is to give you the concepts, the intuitions, and the tools you need to actually learn how deep learning works and how and when to use it. What this book offers... You will learn all about statistics, Probability, Machine Learning, Fundamentals and Deep Learning. All the modules will contain examples and images to help beginner understand. Clear and Easy to Understand Solutions The book is extensively tested by a group of beta readers. The contents provided are simplified as much as possible so that they can serve as starting point for you to more advanced courses. What this book aims to do... This book is written with one goal in mind - to help beginners overcome their initial obstacles to learn Deep Learning from Scratch. A lot of times, newbies tend to feel intimidated by Data Science Models and Coding. The goal of this book is to isolate the different concepts so that beginners can gradually gain competency in the fundamentals and keys concepts before start working on a real project. Beginners in Data Science does not have to be scary or frustrating when you take one step at a time. Ready to start practicing and start learning Deep Learning Fundamentals from Scratch? Click the BUY button now to download this book Topics Covered: Fundamentals of Probability Fundamentals of Statistics Fundamentals of Linear Algebra Introduction to Machine Learning and Deep Learning Fundamentals of Machine Learning Fundamentals of Neural Networks and Deep Learning Deep Learning Parameters and Hyper-parameters Deep Neural Networks Layers Deep Learning Activation Functions Deep Learning Loss Functions Deep Learning Optimization Algorithms Convolutional Neural Network Recurrent Neural Networks LSTM Recursive Neural Networks Bonus Course Conclusion Click the BUY button and download the book now to start learning Deep Learning. ** MONEY BACK GUARANTEE BY AMAZON ** If you aren't satisfied, for more information about the amazon refund service please go to the amazon help platform or contact us by sending an email at contact@aispublishing.net. **GET YOUR COPY NOW, the price will be 24.99\$ soon**

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.

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

Deep Learning-Rajiv Chopra A good book is like a teacher who sits behind the reader and guides him/her accordingly. Deep Learning has been an area of current research.After toiling through the various concepts of Deep Learning, the book slithers around all principles of deep learning. This book highlights in deep the concepts of deep learning so that new projects and researchers can be done.The book serves, both as textbook and as a reference book. Some of the highlights of the book are: Simple Language, Recent Concepts of Machine and Deep Learning explained, MCQ's, Conceptual Short Questions & Answers, Case Studies, Case Tools (like TensorFlow, H2O etc).

Artificial Intelligence Networking: Learning AI Explained for Dummies-Laureano Gallardo 2019-06-14 The term Artificial Intelligence represents a set of software, logic, computer and philosophy disciplines that are intended to make PCs perform functions that were thought to be exclusively human, such as perceiving meaning in written language, or spoken, learn, recognize facial expressions, etc. The field of artificial intelligence has a long history behind it, with many previous advances, such as the recognition of optical characters, which are currently considered as something everyday.Artificial intelligence is an attractive concept for many stakeholders in the business, scientific and governmental world. From the economic point of view, it is very attractive that machines can perform tasks for which human intervention was usually required. An effective artificial intelligence solution can "think" faster and process more information than any human brain. Artificial intelligence also has the potential to extend human capabilities to places where people have difficulty reaching, for example, outer space or remote locations on Earth where human skills such as medical knowledge might prove useful.We could define AI as a technological discipline that aims to design and build machines and programs capable of performing complex tasks with expertise equal to or greater than that of a human being. Bearing in mind that those tasks that seem simpler, for example, speaking, seeing, listening, are precisely the most complex tasks to perform.The computer is a powerful tool because of its ability to represent any discrete system of physical symbols. But the ultimate goal is not for this "science" the computer or its programs, but the study of intelligent behavior and, in particular, human behavior.This is the reason why the AI links with so many other sciences. Thus, in the study of human behavior, natural sciences such as neurophysiology and human sciences such as psychology intervene; in the study of learning processes, pedagogy; in other important aspects, such as the recognition of natural language, linguistics; and in the practical realization of the models, the computer science and in all its surroundings.CONTENTS:INTRODUCTION TO AIHISTORICAL EVOLUTION Prehistory of the AI The birth of the AI The difficult years Current developmentPREMISES 1. Just think about big projects 2. Artificial intelligence is only for technological empires 3. Believe that it is a purposeWHAT IS MACHINE LEARNING?WHAT IS DEEP LEARNING?WHO IS DEVELOPING THIS TECHNOLOGY WORLDWIDE?WHY IS IT IMPORTANT FOR GLOBAL DEVELOPMENT?HOW COULD AI AFFECT THE ECONOMY?WHICH SECTORS WORLDWIDE ARE LEADERS IN THE IMPLEMENTATION OF THIS TYPE OF TECHNOLOGY?ADVANTAGES OF ARTIFICIAL INTELLIGENCEDISADVANTAGES OF ARTIFICIAL INTELLIGENCEWHICH COMPANIES ARE LEADERS IN AI?AI AND BIG DATAIS AI DANGEROUS?AI AND MEDICINESUMMARY AI

Deep Learning and the Game of Go-Max Pumperla 2019-01-06 It's nearly impossible to build a competent Go-playing machine using conventional programming techniques, let alone have it win. By applying advanced AI techniques, in particular deep learning and reinforcement learning, users can train their Go-bot in the rules and tactics of the game. Deep Learning and the Game of Go opens up the world of deep learning and AI by teaching readers to build their own Go-playing machine. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

Pragmatic Machine Learning with Python-Avishek Nag 2020-04-30 An easy-to-understand guide to learn practical Machine Learning techniques with Mathematical foundations KEY FEATURES - A balanced combination of underlying mathematical theories & practical examples with Python code - Coverage of latest topics like multi-label classification, Text Mining, Doc2Vec, Word2Vec, XMeans clustering, unsupervised outlier detection, techniques to deploy ML models in production-grade systems with PMML, etc - Coverage of sufficient & relevant visualization techniques specific to any topic DESCRIPTION This book will be ideal for working professionals who want to learn Machine Learning from scratch. The first chapter will be an introductory chapter to make readers comfortable with the idea of Machine Learning and the required mathematical theories. There will be a balanced combination of underlying mathematical theories corresponding to any Machine Learning topic and its implementation using Python. Most of the implementations will be based on 'scikit-learn,' but other Python libraries like 'Gensim' or 'PyTorch' will also be used for some topics like text analytics or deep learning. The book will be divided into chapters based on primary Machine Learning topics like Classification, Regression, Clustering, Deep Learning, Text Mining, etc. The book will also explain different techniques of putting Machine Learning models into production-grade systems using Big Data or Non-Big Data favors and standards for exporting models. WHAT WILL YOU LEARN - Get familiar with practical concepts of Machine Learning from ground zero - Learn how to deploy Machine Learning models in production - Understand how to do "Data Science Storytelling" - Explore the latest topics in the current industry about Machine Learning WHO THIS BOOK IS FOR This book would be ideal for experienced Software Professionals who are trying to get into the field of Machine Learning. Anyone who wishes to Learn Machine Learning concepts and models in the production lifecycle. TABLE OF CONTENTS 1. Introduction to Machine Learning & Mathematical preliminaries 2. Classification 3. Regression 4. Clustering 5. Deep Learning & Neural Networks 6. Miscellaneous Unsupervised Learning 7. Text Mining 8. Machine Learning models in production 9. Case Studies & Data Science Storytelling