

Neural Network Programming With Java Simple Guide On Neural Networks

Summary Deep Learning and the Game of Go teaches you how to apply the power of deep learning to complex reasoning tasks by building a Go-playing AI. After exposing you to the foundations of machine and deep learning, you'll use Python to build a bot and then teach it the rules of the game. Foreword by Thore Graepel, DeepMind Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The ancient strategy game of Go is an incredible case study for AI. In 2016, a deep learning-based system shocked the Go world by defeating a world champion. Shortly after that, the upgraded AlphaGo Zero crushed the original bot by using deep reinforcement learning to master the game. Now, you can learn those same deep learning techniques by building your own Go bot! About the Book Deep Learning and the Game of Go introduces deep learning by teaching you to build a Go-winning bot. As you progress, you'll apply increasingly complex training techniques and strategies using the Python deep learning library Keras. You'll enjoy watching your bot master the game of Go, and along the way, you'll discover how to apply your new deep learning skills to a wide range of other scenarios! What's inside Build and teach a self-improving game AI Enhance classical game AI systems with deep learning Implement neural networks for deep learning About the Reader All you need are basic Python skills and high school-level math. No deep learning experience required. About the Author Max Pumperla and Kevin Ferguson are experienced deep learning specialists skilled in distributed systems and data science. Together, Max and Kevin built the open source bot BetaGo. Table of Contents PART 1 - FOUNDATIONS Toward deep learning: a machine-learning introduction Go as a machine-learning problem Implementing your first Go bot PART 2 - MACHINE LEARNING AND GAME AI Playing games with tree search Getting started with neural networks Designing a neural network for Go data Learning from data: a deep-learning bot Deploying bots in the wild Learning by practice: reinforcement learning Reinforcement learning with policy gradients Reinforcement learning with value methods Reinforcement learning with actor-critic methods PART 3 - GREATER THAN THE SUM OF ITS PARTS AlphaGo: Bringing it all together AlphaGo Zero: Integrating tree search with reinforcement learning Encog is an advanced neural network and bot programming framework. This book focuses on using Encog to create a variety of neural network architectures using the Java programming language. Neural network architectures such as feedforward/perceptrons, Hopfield, Elman, Jordan, Radial Basis Function, and Self Organizing maps are all demonstrated. This book also shows how to use Encog to train neural networks using a variety of means. Several propagation techniques, such as back propagation, resilient propagation (RPROP) and the Manhattan update rule are discussed. Additionally, training with a genetic algorithm and simulated annealing is discussed as well. You will also see how to enhance training using techniques such as pruning, hybrid training, Real world examples tie the book together. Pattern recognition applications such as OCR, image and text recognition will be introduced. You will see how to apply time series and forecasting and how to financial markets. All of the Encog neural network components will be demonstrated to show how to use them in your own neural network applications.

This text serves as a cookbook for neural network solutions to practical problems using C++. It will enable those with moderate programming experience to select a neural network model appropriate to solving a particular problem, and to produce a working program implementing that network. The book provides guidance along the entire problem-solving path, including designing the training set, preprocessing variables, training and validating the network, and evaluating its performance. Though the book is not intended as a general course in neural networks, no background in neural works is assumed and all models are presented from the ground up. The principle focus of the book is the three layer feedforward network, for more than a decade as the workhorse of professional arsenals. Other network models with strong performance records are also included. Bound in the book is an IBM diskette that includes the source code for all programs in the book. Much of this code can be easily adapted to C compilers. In addition, the operation of all programs is thoroughly discussed both in the text and in the comments within the code to facilitate translation to other languages.

Create and unleash the power of neural networks by implementing professional, clean, and clear Java code About This Book* Learn to build amazing projects using neural networks including forecasting the weather and pattern recognition* Explore the Java multi-platform feature to run your personal neural networks everywhere* This step-by-step guide will help you solve real-world problems and links neural network theory to their application Who This Book Is For This book is for Java developers who want to know how to develop smarter applications using the power of neural networks. Those who deal with a lot of complex data and want to use it efficiently in their day-to-day apps will find this book quite useful. Some basic experience with statistical computations is expected. What You Will Learn* Develop an understanding of neural networks and how they can be fitted* Explore the learning process of neural networks* Build neural network applications with Java using hands-on examples* Discover the power of neural network's unsupervised learning process to extract the intrinsic knowledge hidden behind the data* Apply the code generated in practical examples, including weather forecasting and pattern recognition* Understand how to make the best choice of learning parameters to ensure you have a more effective application* Select and split data sets into training, test, and validation, and explore validation strategies In Detail Want to discover the current state-of-art in the field of neural networks that will let you understand and design new strategies to apply to more complex problems? This book takes you on a complete walkthrough of the process of developing basic to advanced practical examples based on neural networks with Java, giving you everything you need to stand out. You will first learn the basics of neural networks and their process of learning. We then focus on what Perceptrons are and their features. Next, you will implement self-organizing maps using practical examples. Further on, you will learn about some of the applications that are presented in this book such as weather forecasting, disease diagnosis, customer profiling, generalization, extreme machine learning, and characters recognition (OCR). Finally, you will learn methods to optimize and adapt neural networks in real time. All the examples generated in the book are provided in the form of illustrative source code, which merges object-oriented programming (OOP) concepts and neural network features to enhance your learning experience.

Use Java and Deeplearning4j to build robust, scalable, and highly accurate AI models from scratch Key Features Install and configure Deeplearning4j to implement deep learning models from scratch Explore recipes for developing, training, and fine-tuning your neural network models in Java Model neural networks using datasets containing images, text, and time-series data Book Description Java is one of the most widely used programming languages in the world. With this book, you will see how to perform deep learning using Deeplearning4j (DL4J) - the most popular Java library for training neural networks efficiently. This book starts by showing you how to install and configure Java and DL4J on your system. You will then gain insights into deep learning basics and use your knowledge to create a deep neural network for binary classification from scratch. As you progress, you will discover how to build a convolutional neural network (CNN) in DL4J, and understand how to construct numeric vectors from text. This deep learning book will also guide you through performing anomaly detection on unsupervised data and help you set up neural networks in distributed systems effectively. In addition to this, you will learn how to import models from Keras and change the configuration in a pre-trained DL4J model. Finally, you will explore benchmarking in DL4J and optimize neural networks for optimal results. By the end of this book, you will have a clear understanding of how you can use DL4J to build robust deep learning applications in Java. What you will learn Perform data normalization and wrangling using DL4J Build deep neural networks using DL4J Implement CNNs to solve image classification problems Train autoencoders to solve anomaly detection problems using DL4J Perform benchmarking and optimization to improve your model's performance Implement reinforcement learning for real-world use cases using RL4J Leverage the capabilities of DL4J in distributed systems Who this book is for If you are a data scientist, machine learning developer, or a deep learning enthusiast who wants to implement deep learning models in Java, this book is for you. Basic understanding of

Java programming as well as some experience with machine learning and neural networks is required to get the most out of this book. Build and run intelligent applications by leveraging key Java machine learning libraries About This Book Develop a sound strategy to solve predictive modelling problems using the most popular machine learning Java libraries. Explore a broad variety of data processing, machine learning, and natural language processing through diagrams, source code, and real-world applications This step-by-step guide will help you solve real-world problems and links neural network theory to their application Who This Book Is For This course is intended for data scientists and Java developers who want to dive into the exciting world of deep learning. It will get you up and running quickly and provide you with the skills you need to successfully create, customize, and deploy machine learning applications in real life. What You Will Learn Get a practical deep dive into machine learning and deep learning algorithms Explore neural networks using some of the most popular Deep Learning frameworks Dive into Deep Belief Nets and Stacked Denoising Autoencoders algorithms Apply machine learning to fraud, anomaly, and outlier detection Experiment with deep learning concepts, algorithms, and the toolbox for deep learning Select and split data sets into training, test, and validation, and explore validation strategies Apply the code generated in practical examples, including weather forecasting and pattern recognition In Detail Machine learning applications are everywhere, from self-driving cars, spam detection, document search, and trading strategies, to speech recognition Starting with an introduction to basic machine learning algorithms, this course takes you further into this vital world of stunning predictive insights and remarkable machine intelligence. This course helps you solve challenging problems in image processing, speech recognition, language modeling. You will discover how to detect anomalies and fraud, and ways to perform activity recognition, image recognition, and text. You will also work with examples such as weather forecasting, disease diagnosis, customer profiling, generalization, extreme machine learning and more. By the end of this course, you will have all the knowledge you need to perform deep learning on your system with varying complexity levels, to apply them to your daily work. The course provides you with highly practical content explaining deep learning with Java, from the following Packt books: Java Deep Learning Essentials Machine Learning in Java Neural Network Programming with Java, Second Edition Style and approach This course aims to create a smooth learning path that will teach you how to effectively use deep learning with Java with other de facto components to get the most out of it. Through this comprehensive course, you'll learn the basics of predictive modelling and progress to solve real-world problems and links neural network theory to their application "We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document." —Soumith Chintala, co-creator of PyTorch Key Features Written by PyTorch's creator and key contributors Develop deep learning models in a familiar Pythonic way Use PyTorch to build an image classifier for cancer detection Diagnose problems with your neural network and improve training with data augmentation Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About The Book 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. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you'll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks. What You Will Learn Understanding deep learning data structures such as tensors and neural networks Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Methods for training networks with limited inputs Sifting through unreliable results to diagnose and fix problems in your neural network Improve your results with augmented data, better model architecture, and fine tuning This Book Is Written For For Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required. About The Authors Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and 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

Design, build, and deploy your own machine learning applications by leveraging key Java machine learning libraries About This Book- Develop a sound strategy to solve predictive modelling problems using the most popular machine learning Java libraries- Explore a broad variety of data processing, machine learning, and natural language processing through diagrams, source code, and real-world applications- Packed with practical advice and tips to help you get to grips with applied machine learning Who This Book Is For If you want to learn how to use Java's machine learning libraries to gain insight from your data, this book is for you. It will get you up and running quickly and provide you with the skills you need to successfully create, customize, and deploy machine learning applications in real life. You should be familiar with Java programming and data mining concepts to make the most of this book, but no prior experience with data mining packages is necessary. What You Will Learn- Understand the basic steps of applied machine learning and how to differentiate among various machine learning approaches- Discover key Java machine learning libraries, what each library brings to the table, and what kind of problems each are able to solve- Learn how to implement classification, regression, and clustering- Develop a sustainable strategy for customer retention by predicting likely churn candidates- Build a scalable recommendation engine with Apache Mahout- Apply machine learning to fraud, anomaly, and outlier detection- Experiment with deep learning concepts, algorithms, and the toolbox for deep learning- Write your own activity recognition model for eHealth applications using mobile sensors In Detail As the amount of data continues to grow at an almost incomprehensible rate, being able to understand and process data is becoming a key differentiator for competitive organizations. Machine learning applications are everywhere, from self-driving cars, spam detection, document search, and trading strategies, to speech recognition. This makes machine learning well-suited to the present-day era of Big Data and Data Science. The main challenge is how to transform data into actionable knowledge. Machine Learning in Java will provide you with the techniques and tools you need to quickly gain insight from complex data. You will start by learning how to apply machine learning methods to a variety of common tasks including classification, prediction, forecasting, market basket analysis, and clustering. Moving on, you will discover how to detect anomalies and fraud, and ways to perform activity recognition, image recognition, and text analysis. By the end of the book, you will explore related web resources and technologies that will help you take your learning to the next level. By applying the most effective machine learning methods to real-world problems, you will gain hands-on experience that will transform the way you think about data. Style and approach This is a practical tutorial that uses hands-on examples to step through some real-world applications of machine learning. Without shying away from the technical details, you will explore machine learning with Java libraries using clear and practical examples. You will explore how to prepare data for analysis, choose a machine learning method, and measure the success of the process.

This book shows you how to build real-time image processing systems all the way through to house automation. Find out how you can develop a system based on small 32-bit ARM processors that gives you complete control through voice commands. Real-time image processing systems are utilized in a wide variety of applications, such as in traffic monitoring systems, medical image processing, and biometric security systems. In Real-Time IoT Imaging with Deep Neural Networks, you will learn how to make use of the best DNN models to detect object in images using Java and a wrapper for OpenCV. Take a closer look at how Java scripting works on the Raspberry Pi while preparing your Visual Studio code for remote programming. You will also gain insights on image and video scripting. Author Nicolas Modrzyk shows you how to use the Rhasspy voice platform to add a powerful voice assistant and completely run and control your Raspberry Pi from your computer. To get your voice intents for house automation ready, you will explore how Java connects to the MQTT and handles parametrized Rhasspy voice commands. With your voice-controlled system ready for operation, you will be able to perform simple tasks such as detecting cats, people, and coffee pots in your selected environment. Privacy and freedom are essential, so priority is given to using open source software and an on-device voice environment where you have full control of your data and video streams. Your voice commands are your own—and just your own. With recent advancements in the Internet of Things and machine learning, cutting edge image processing systems provide complete process automation. This practical book teaches you to build such a system, giving you complete control with minimal effort. What You Will Learn: Show mastery by creating OpenCV filters Execute a YOLO DNN model for image detection Apply the best Java scripting on Raspberry Pi 4 Prepare your setup for real-time remote programming Use the Rhasspy voice platform for handling voice commands and enhancing your house automation setup Who This Book Is For:Engineers, and Hobbyists wanting to use their favorite JVM to run Object Detection and Networks on a Raspberry Pi

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

Neural networks are a computing paradigm that is finding increasing attention among computer scientists. In this book, theoretical laws and models previously scattered in the literature are brought together into a general theory of artificial neural nets. Always with a view to biology and starting with the simplest nets, it is shown how the properties of models change when more general computing elements and net topologies are introduced. Each chapter contains examples, numerous illustrations, and a bibliography. The book is aimed at readers who seek an overview of the field or who wish to deepen their knowledge. It is suitable as a basis for university courses in neurocomputing.

Introduction to Neural Networks in Java, Second Edition, introduces the Java programmer to the world of Neural Networks and Artificial Intelligence. Neural network architectures such as the feedforward, Hopfield, and Self Organizing Map networks are discussed. Training techniques such as Backpropagation, Genetic Algorithms and Simulated Annealing are also introduced. Practical examples are given for each neural network. Examples include the Traveling Salesman problem, handwriting recognition, financial prediction, game strategy, learning mathematical functions and special application to Internet bots. All Java source code can be downloaded online.

An unparalleled opportunity to learn about an exciting new technology that is revolutionizing network and Internet content delivery Network Query Language (NQL) is a revolutionary new scripting language that makes it astonishingly quick and easy to aggregate, analyze, interpret, and redistribute information via networks. Described as the "first language of the content engineering era," NQL allows programmers to develop bots, intelligent agents, middleware, and sophisticated Web applications in a small fraction of the time it would ordinarily take. This book offers developers and network administrators an unparalleled opportunity to learn about this exciting new technology-what it is, how it works, and how to use and make the most of its many features-from the man who invented it. CD-ROM contains a 60-day-timeout version of NQL in both Windows and Java (Linux and Macintosh) formats.

Use Java and Deeplearning4j to build robust, scalable, and highly accurate AI models from scratch Key Features Install and configure Deeplearning4j to implement deep learning models from scratch Explore recipes for developing, training, and fine-tuning your neural network models in Java Model neural networks using datasets containing images, text, and time-series data Book Description Java is one of the most widely used programming languages in the world. With this

book, you will see how to perform deep learning using Deeplearning4j (DL4J) – the most popular Java library for training neural networks efficiently. This book starts by showing you how to install and configure Java and DL4J on your system. You will then gain insights into deep learning basics and use your knowledge to create a deep neural network for binary classification from scratch. As you progress, you will discover how to build a convolutional neural network (CNN) in DL4J, and understand how to construct numeric vectors from text. This deep learning book will also guide you through performing anomaly detection on unsupervised data and help you set up neural networks in distributed systems effectively. In addition to this, you will learn how to import models from Keras and change the configuration in a pre-trained DL4J model. Finally, you will explore benchmarking in DL4J and optimize neural networks for optimal results. By the end of this book, you will have a clear understanding of how you can use DL4J to build robust deep learning applications in Java. What you will learn

- Perform data normalization and wrangling using DL4J
- Build deep neural networks using DL4J
- Implement CNNs to solve image classification problems
- Train autoencoders to solve anomaly detection problems using DL4J
- Perform benchmarking and optimization to improve your model's performance
- Implement reinforcement learning for real-world use cases using RL4J
- Leverage the capabilities of DL4J in distributed systems

Who this book is for If you are a data scientist, machine learning developer, or a deep learning enthusiast who wants to implement deep learning models in Java, this book is for you. Basic understanding of Java programming as well as some experience with machine learning and neural networks is required to get the most out of this book.

Create and unleash the power of neural networks by implementing professional Java code

About This Book

- Learn to build amazing projects using neural networks including forecasting the weather and pattern recognition
- Explore the Java multi-platform feature to run your personal neural networks everywhere
- This step-by-step guide will help you solve real-world problems and links neural network theory to their application

Who This Book Is For

This book is for Java developers with basic Java programming knowledge. No previous knowledge of neural networks is required as this book covers the concepts from scratch.

What You Will Learn

- Get to grips with the basics of neural networks and what they are used for
- Develop neural networks using hands-on examples
- Explore and code the most widely-used learning algorithms to make your neural network learn from most types of data
- Discover the power of neural network's unsupervised learning process to extract the intrinsic knowledge hidden behind the data
- Apply the code generated in practical examples, including weather forecasting and pattern recognition
- Understand how to make the best choice of learning parameters to ensure you have a more effective application
- Select and split data sets into training, test, and validation, and explore validation strategies
- Discover how to improve and optimize your neural network

In Detail

Vast quantities of data are produced every second. In this context, neural networks become a powerful technique to extract useful knowledge from large amounts of raw, seemingly unrelated data. One of the most preferred languages for neural network programming is Java as it is easier to write code using it, and most of the most popular neural network packages around already exist for Java. This makes it a versatile programming language for neural networks. This book gives you a complete walkthrough of the process of developing basic to advanced practical examples based on neural networks with Java. You will first learn the basics of neural networks and their process of learning. We then focus on what Perceptrons are and their features. Next, you will implement self-organizing maps using the concepts you've learned. Furthermore, you will learn about some of the applications that are presented in this book such as weather forecasting, disease diagnosis, customer profiling, and characters recognition (OCR). Finally, you will learn methods to optimize and adapt neural networks in real time. All the examples generated in the book are provided in the form of illustrative source code, which merges object-oriented programming (OOP) concepts and neural network features to enhance your learning experience.

Style and approach

This book adopts a step-by-step approach to neural network development and provides many hands-on examples using Java programming. Each neural network concept is explored through real-world problems and is delivered in an easy-to-comprehend manner.

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

Though mathematical ideas underpin the study of neural networks, the author presents the fundamentals without the full mathematical apparatus. All aspects of the field are tackled, including artificial neurons as models of their real counterparts; the geometry of network action in pattern space; gradient descent methods, including back-propagation; associative memory and Hopfield nets; and self-organization and feature maps. The traditionally difficult topic of adaptive resonance theory is clarified within a hierarchical description of its operation. The book also includes several real-world examples to provide a concrete focus. This should enhance its appeal to those involved in the design, construction and management of networks in commercial environments and who wish to improve their understanding of network simulator packages. As a comprehensive and highly accessible introduction to one of the most important topics in cognitive and computer science, this volume should interest a wide range of readers, both students and professionals, in cognitive science, psychology, computer science and electrical engineering.

The idea of simulating the brain was the goal of many pioneering works in Artificial Intelligence. The brain has been seen

as a neural network, or a set of nodes, or neurons, connected by communication lines. Currently, there has been increasing interest in the use of neural network models. This book contains chapters on basic concepts of artificial neural networks, recent connectionist architectures and several successful applications in various fields of knowledge, from assisted speech therapy to remote sensing of hydrological parameters, from fabric defect classification to application in civil engineering. This is a current book on Artificial Neural Networks and Applications, bringing recent advances in the area to the reader interested in this always-evolving machine learning technique.

In addition to showing the programmer how to construct Neural Networks, the book discusses the Java Object Oriented Neural Engine (JOONE), a free open source Java neural engine. (Computers)

Summary Deep learning has transformed the fields of computer vision, image processing, and natural language applications. Thanks to TensorFlow.js, now JavaScript developers can build deep learning apps without relying on Python or R. Deep Learning with JavaScript shows developers how they can bring DL technology to the web. Written by the main authors of the TensorFlow library, this new book provides fascinating use cases and in-depth instruction for deep learning apps in JavaScript in your browser or on Node. Foreword by Nikhil Thorat and Daniel Smilkov. About the technology Running deep learning applications in the browser or on Node-based backends opens up exciting possibilities for smart web applications. With the TensorFlow.js library, you build and train deep learning models with JavaScript. Offering uncompromising production-quality scalability, modularity, and responsiveness, TensorFlow.js really shines for its portability. Its models run anywhere JavaScript runs, pushing ML farther up the application stack. About the book In Deep Learning with JavaScript, you'll learn to use TensorFlow.js to build deep learning models that run directly in the browser. This fast-paced book, written by Google engineers, is practical, engaging, and easy to follow. Through diverse examples featuring text analysis, speech processing, image recognition, and self-learning game AI, you'll master all the basics of deep learning and explore advanced concepts, like retraining existing models for transfer learning and image generation. What's inside - Image and language processing in the browser - Tuning ML models with client-side data - Text and image creation with generative deep learning - Source code samples to test and modify About the reader For JavaScript programmers interested in deep learning. About the author Shuang Cai, Stanley Bileschi and Eric D. Nielsen are software engineers with experience on the Google Brain team, and were crucial to the development of the high-level API of TensorFlow.js. This book is based in part on the classic, Deep Learning with Python by François Chollet. TOC: PART 1 - MOTIVATION AND BASIC CONCEPTS 1 • Deep learning and JavaScript PART 2 - A GENTLE INTRODUCTION TO TENSORFLOW.JS 2 • Getting started: Simple linear regression in TensorFlow.js 3 • Adding nonlinearity: Beyond weighted sums 4 • Recognizing images and sounds using convnets 5 • Transfer learning: Reusing pretrained neural networks PART 3 - ADVANCED DEEP LEARNING WITH TENSORFLOW.JS 6 • Working with data 7 • Visualizing data and models 8 • Underfitting, overfitting, and the universal workflow of machine learning 9 • Deep learning for sequences and text 10 • Generative deep learning 11 • Basics of deep reinforcement learning PART 4 - SUMMARY AND CLOSING WORDS 12 • Testing, optimizing, and deploying models 13 • Summary, conclusions, and beyond

Develop neural network applications using the Java environment. After learning the rules involved in neural network processing, this second edition shows you how to manually process your first neural network example. The book covers the internals of front and back propagation and helps you understand the main principles of neural network processing. You also will learn how to prepare the data to be used in neural network development and you will be able to suggest various techniques of data preparation for many unconventional tasks. This book discusses the practical aspects of using Java for neural network processing. You will know how to use the Encog Java framework for processing large-scale neural network applications. Also covered is the use of neural networks for approximation of non-continuous functions. In addition to using neural networks for regression, this second edition shows you how to use neural networks for computer vision. It focuses on image recognition such as the classification of handwritten digits, input data preparation and conversion, and building the conversion program. And you will learn about topics related to the classification of handwritten digits such as network architecture, program code, programming logic, and execution. The step-by-step approach taken in the book includes plenty of examples, diagrams, and screenshots to help you grasp the concepts quickly and easily. What You Will Learn Use Java for the development of neural network applications Prepare data for many different tasks Carry out some unusual neural network processing Use a neural network to process non-continuous functions Develop a program that recognizes handwritten digits Who This Book Is For Intermediate machine learning and deep learning developers who are interested in switching to Java

A step-by-step gentle journey through the mathematics of neural networks, and making your own using the Python computer language. Neural networks are a key element of deep learning and artificial intelligence, which today is capable of some truly impressive feats. Yet too few really understand how neural networks actually work. This guide will take you on a fun and unhurried journey, starting from very simple ideas, and gradually building up an understanding of how neural networks work. You won't need any mathematics beyond secondary school, and an accessible introduction to calculus is also included. The ambition of this guide is to make neural networks as accessible as possible to as many readers as possible - there are enough texts for advanced readers already! You'll learn to code in Python and make your own neural network, teaching it to recognise human handwritten numbers, and performing as well as professionally developed networks. Part 1 is about ideas. We introduce the mathematical ideas underlying the neural networks, gently with lots of illustrations and examples. Part 2 is practical. We introduce the popular and easy to learn Python programming language, and gradually builds up a neural network which can learn to recognise human handwritten numbers, easily getting it to perform as well as networks made by professionals. Part 3 extends these ideas further. We push the performance of our neural network to an industry leading 98% using only simple ideas and code, test the

network on your own handwriting, take a privileged peek inside the mysterious mind of a neural network, and even get it all working on a Raspberry Pi. All the code in this has been tested to work on a Raspberry Pi Zero.

This book is an exploration of neural networks and how to implement them in Java. First, the reader is guided so as to understand what neural networks are. You will learn how they operate. The process of learning in neural networks is very important. This is the concept which makes neural networks behave in the same manner as the brain of human beings. This process is discussed in this book. You are also guided on how to implement this in Java. The Java lego robots are very common in the field of artificial intelligence. This book guides you on how to implement these in Java. Recurrent neural networks, which are believed to have memory, are discussed in detail. These work in such a way that the value will be calculated based on the value obtained in the previous step. You will learn how to implement such a network in Java. Convolutional neural networks are also explored in detail. You will learn how these work as well as how to implement them in Java. The following topics are discussed in this book: -Understanding Neural Networks -Learning in Neural Networks -Java Lego Robots Neural Network -Convolutional Neural Networks -Recurrent Neural Networks

This book will take you through the process of efficiently training deep neural networks in Java for Computer Vision-related tasks. You will build real-world applications ranging from simple Java handwritten digit recognition models to real-time autonomous car driving systems and face recognition models using the popular Java-based libraries.

Books on computation in the marketplace tend to discuss the topics within specific fields. Many computational algorithms, however, share common roots. Great advantages emerge if numerical methodologies break the boundaries and find their uses across disciplines. Interdisciplinary Computing In Java Programming Language introduces readers of different backgrounds to the beauty of the selected algorithms. Serious quantitative researchers, writing customized codes for computation, enjoy cracking source codes as opposed to the black-box approach. Most C and Fortran programs, despite being slightly faster in program execution, lack built-in support for plotting and graphical user interface. This book selects Java as the platform where source codes are developed and applications are run, helping readers/users best appreciate the fun of computation. Interdisciplinary Computing In Java Programming Language is designed to meet the needs of a professional audience composed of practitioners and researchers in science and technology. This book is also suitable for senior undergraduate and graduate-level students in computer science, as a secondary text.

Learn practical uses for some of the hottest tech applications trending among technology professionals We are living in an era of digital revolution. On the horizon, many emerging digital technologies are being developed at a breathtaking speed. Whether we like it or not, whether we are ready or not, digital technologies are going to penetrate more and more, deeper and deeper, into every aspect of our lives. This is going to fundamentally change how we live, how we work, and how we socialize. Java, as a modern high-level programming language, is an excellent tool for helping us to learn these digital technologies, as well as to develop digital applications, such as IoT, AI, Cybersecurity, Blockchain and more.

Practical Java Programming uses Java as a tool to help you learn these new digital technologies and to be better prepared for the future changes. Gives you a brief overview for getting started with Java Programming Dives into how you can apply your new knowledge to some of the biggest trending applications today Helps you understand how to program Java to interact with operating systems, networking, and mobile applications Shows you how Java can be used in trending tech applications such as IoT (Internet of Things), AI (Artificial Intelligence), Cybersecurity, and Blockchain Get ready to find out firsthand how Java can be used for connected home devices, healthcare, the cloud, and all the hottest tech applications.

Summary Deep Learning for Search teaches you how to improve the effectiveness of your search by implementing neural network-based techniques. By the time you're finished with the book, you'll be ready to build amazing search engines that deliver the results your users need and that get better as time goes on! Foreword by Chris Mattmann. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning handles the toughest search challenges, including imprecise search terms, badly indexed data, and retrieving images with minimal metadata. And with modern tools like DL4J and TensorFlow, you can apply powerful DL techniques without a deep background in data science or natural language processing (NLP). This book will show you how. About the Book Deep Learning for Search teaches you to improve your search results with neural networks. You'll review how DL relates to search basics like indexing and ranking. Then, you'll walk through in-depth examples to upgrade your search with DL techniques using Apache Lucene and Deeplearning4j. As the book progresses, you'll explore advanced topics like searching through images, translating user queries, and designing search engines that improve as they learn! What's inside Accurate and relevant rankings Searching across languages Content-based image search Search with recommendations About the Reader For developers comfortable with Java or a similar language and search basics. No experience with deep learning or NLP needed. About the Author Tommaso Teofili is a software engineer with a passion for open source and machine learning. As a member of the Apache Software Foundation, he contributes to a number of open source projects, ranging from topics like information retrieval (such as Lucene and Solr) to natural language processing and machine translation (including OpenNLP, Joshua, and UIMA). He currently works at Adobe, developing search and indexing infrastructure components, and researching the areas of natural language processing, information retrieval, and deep learning. He has presented search and machine learning talks at conferences including BerlinBuzzwords, International Conference on Computational Science, ApacheCon, EclipseCon, and others. You can find him on Twitter at @tteofili. Table of Contents PART 1 - SEARCH MEETS DEEP LEARNING Neural search Generating synonyms PART 2 - THROWING NEURAL NETS AT A SEARCH ENGINE From plain retrieval to text generation More-sensitive query suggestions Ranking search results with word embeddings Document embeddings for rankings and recommendations PART 3 - ONE STEP BEYOND Searching across languages Content-based image search A peek at performance

The study of nonlinear dynamical systems has advanced tremendously in the last 20 years, making a big impact on science and technology. This book provides all the techniques and methods used in nonlinear dynamics. The concepts and underlying mathematics are discussed in detail. The numerical and symbolic methods are implemented in C++, SymbolicC++ and Java. Object-oriented techniques are also applied. The book contains more than 150 ready-to-run programs. The text has also been designed for a one-year course at both the junior and senior levels in nonlinear dynamics. The topics discussed in the book are part of e-learning and distance learning courses conducted by the International School for Scientific Computing, University of Johannesburg.

This resource introduces the C# programmer to the world of Neural Networks and Artificial Intelligence. Training techniques, such as backpropagation, genetic algorithms, and simulated annealing are also introduced.

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

Build and run intelligent applications by leveraging key Java machine learning libraries About This Book* Develop a sound strategy to solve predictive modelling problems using the most popular machine learning Java libraries.* Explore a broad variety of data processing, machine learning, and natural language processing through diagrams, source code, and real-world applications* This step-by-step guide will help you solve real-world problems and links neural network theory to their application Who This Book Is For This course is intended for data scientists and Java developers who want to dive into the exciting world of deep learning. It will get you up and running quickly and provide you with the skills you need to successfully create, customize, and deploy machine learning applications in real life. What You Will Learn* Get a practical deep dive into machine learning and deep learning algorithms* Explore neural networks using some of the most popular Deep Learning frameworks* Dive into Deep Belief Nets and Stacked Denoising Autoencoders algorithms* Apply machine learning to fraud, anomaly, and outlier detection* Experiment with deep learning concepts, algorithms, and the toolbox for deep learning* Select and split data sets into training, test, and validation, and explore validation strategies* Apply the code generated in practical examples, including weather forecasting and pattern recognition In Detail Machine learning applications are everywhere, from self-driving cars, spam detection, document search, and trading strategies, to speech recognition Starting with an introduction to basic machine learning algorithms, this course takes you further into this vital world of stunning predictive insights and remarkable machine intelligence. This course helps you solve challenging problems in image processing, speech recognition, language modeling. You will discover how to detect anomalies and fraud, and ways to perform activity recognition, image recognition, and text. You will also work with examples such as weather forecasting, disease diagnosis, customer profiling, generalization, extreme machine learning and more. By the end of this course, you will have all the knowledge you need to perform deep learning on your system with varying complexity levels, to apply them to your daily work. The course provides you with highly practical content explaining deep learning with Java, from the following Packt books: 1. Java Deep Learning Essentials 2. Machine Learning in Java 3. Neural Network Programming with Java, Second Edition Style and approach This course aims to create a smooth learning path that will teach you how to effectively use deep learning with Java with other de facto components to get the most out of it. Through this comprehensive course, you'll learn the basics of predictive modelling and progress to solve real-world problems and links neural network theory to their application Use Java to develop neural network applications in this practical book. After learning the rules involved in neural network processing, you will manually process the first neural network example. This covers the internals of front and back propagation, and facilitates the understanding of the main principles of neural network processing. Artificial Neural Networks with Java also teaches you how to prepare the data to be used in neural network development and suggests various techniques of data preparation for many unconventional tasks. The next big topic discussed in the book is using Java for neural network processing. You will use the Encog Java framework and discover how to do rapid development with Encog, allowing you to create large-scale neural network applications. The book also discusses the inability of neural networks to approximate complex non-continuous functions, and it introduces the micro-batch method that solves this issue. The step-by-step approach includes plenty of examples, diagrams, and screen shots to help you grasp the concepts quickly and easily. What You Will Learn Prepare your data for many different tasks Carry out some unusual neural network tasks Create neural network to process non-continuous functions Select and improve the development model Who This Book Is For Intermediate machine learning and deep learning developers who are interested in switching to Java.

Beginning where our introductory neural network programming book left off, this book introduces you to Encog. Encog allows you to focus less on the actual implementation of neural networks and focus on how to use them. Encog is an advanced neural network programming framework that allows you to create a variety of neural network architectures using the Java programming language. Neural network architectures such as feedforward/perceptrons, Hopfield, Elman, Jordan, Radial Basis Function, and Self Organizing maps are all demonstrated. This book also shows how to use Encog to train neural networks using a variety of means. Several propagation techniques, such as back propagation, resilient propagation (RPROP) and the Manhattan update rule are discussed. Additionally, training with a genetic algorithm and simulated annealing is discussed as well. You will also see how to enhance training using techniques such as pruning and hybrid training.

You will build full-fledged, deep learning applications with Java and different open-source libraries. Master numerical computing, deep learning, and the latest Java programming features to carry out complex advanced tasks. This book is filled with best practices/tips after every project to help you optimize your deep learning models with ease.

Build smarter programs with the power of neural networks and the simplicity of Python About This Book* Make your roots stronger in neural networks by this concept-rich yet highly practical guide; from single layer to multiple layers with the help of Python* Through this book, you will develop a strong background in neural networks, regardless of your level of previous knowledge in this subject* You will be able to implement solutions from scratch, so the whole process on foundations of neural network solution design will be paced by you Who This Book Is For This book is designed for novices as well as intermediate Python developers who have a statistical background and want to work with neural networks to get better results from complex data. It also contains enough food for thought for those who want to improve their skills in machine learning and deep learning. What You Will Learn* See the latest innovations in the field* Become fluent in Python to develop neural networks solutions capable of solving complex and interesting tasks* Implement neural networks step-by-step* Solve your complex computational problems with the aid of neural networks and Python* The reader will be able to set up his/her neural network with ease, according to the objective he/she wants to apply.* The reader will be able to design time series based models using RNNs in Python.* Will be able to design high level solutions with CNNs in Python In Detail If you wish to solve your complex computational problem efficiently, neural networks come to the rescue. This book will teach you how to ace neural networks and solve your computational problems with Python-right

from predicting to self-learning models-with ease. We start off with neural network design, then you'll build a solid foundational knowledge of how a neural network learns from data, and the principles behind it. This book cover various types of neural networks including recurrent neural networks and convoluted neural networks. You will not only learn how to train neural networks, but also see a generalization of these networks. With the help of practical examples and real-world use cases, you will learn to implement these neural networks in your applications.

This book explores the intuitive appeal of neural networks and the genetic algorithm in finance. It demonstrates how neural networks used in combination with evolutionary computation outperform classical econometric methods for accuracy in forecasting, classification and dimensionality reduction. McNelis utilizes a variety of examples, from forecasting automobile production and corporate bond spread, to inflation and deflation processes in Hong Kong and Japan, to credit card default in Germany to bank failures in Texas, to cap-floor volatilities in New York and Hong Kong. * Offers a balanced, critical review of the neural network methods and genetic algorithms used in finance * Includes numerous examples and applications * Numerical illustrations use MATLAB code and the book is accompanied by a website

Simulation in NSL - Modeling in NSL - Schematic Capture System - User Interface and Graphical Windows - The Modeling Language NSLM - The Scripting Language NSLS - Adaptive Resonance Theory - Depth Perception - Retina - Receptive Fields - The Associative Search Network: Landmark Learning and Hill Climbing - A Model of Primate Visual-Motor Conditional Learning - The Modular Design of the Oculomotor System in Monkeys - Crowley-Arbib Saccade Model - A Cerebellar Model of Sensorimotor Adaptation - Learning to Detour - Face Recognition by Dynamic Link Matching - Appendix I : NSLM Methods - NSLJ Extensions - NSLC Extensions - NSLJ and NSLC Differences - NSLJ and NSLC Installation Instructions.

Harness the hidden power of Java to build network-enabled applications with lower network traffic and faster processes

About This Book Learn to deliver superior server-to-server communication through the networking channels Gain expertise of the networking features of your own applications to support various network architectures such as client/server and peer-to-peer Explore the issues that impact scalability, affect security, and allow applications to work in a heterogeneous environment Who This Book Is For Learning Network Programming with Java is oriented to developers who wish to use network technologies to enhance the utility of their applications. You should have a working knowledge of Java and an interest in learning the latest in network programming techniques using Java. No prior experience with network development or special software beyond the Java SDK is needed. Upon completion of the book, beginner and experienced developers will be able to use Java to access resources across a network and the Internet. What You Will Learn Connect to other applications using sockets Use channels and buffers to enhance communication between applications Access network services and develop client/server applications Explore the critical elements of peer-to-peer applications and current technologies available Use UDP to perform multicasting Address scalability through the use of core and advanced threading techniques Incorporate techniques into an application to make it more secure Configure and address interoperability issues to enable your applications to work in a heterogeneous environment In Detail Network-aware applications are becoming more prevalent and play an ever-increasing role in the world today. Connecting and using an Internet-based service is a frequent requirement for many applications. Java provides numerous classes that have evolved over the years to meet evolving network needs. These range from low-level socket and IP-based approaches to those encapsulated in software services. This book explores how Java supports networks, starting with the basics and then advancing to more complex topics. An overview of each relevant network technology is presented followed by detailed examples of how to use Java to support these technologies. We start with the basics of networking and then explore how Java supports the development of client/server and peer-to-peer applications. The NIO packages are examined as well as multitasking and how network applications can address practical issues such as security. A discussion on networking concepts will put many network issues into perspective and let you focus on the appropriate technology for the problem at hand. The examples used will provide a good starting point to develop similar capabilities for many of your network needs. Style and approach Each network technology's terms and concepts are introduced first. This is followed up with code examples to explain these technologies. Many of the examples are supplemented with alternate Java 8 solutions when appropriate. Knowledge of Java 8 is not necessary but these examples will help you better understand the power of Java 8.

Create and unleash the power of neural networks by implementing C# and .Net code Key Features Get a strong foundation of neural networks with access to various machine learning and deep learning libraries Real-world case studies illustrating various neural network techniques and architectures used by practitioners Cutting-edge coverage of Deep Networks, optimization algorithms, convolutional networks, autoencoders and many more Book Description Neural networks have made a surprise comeback in the last few years and have brought tremendous innovation in the world of artificial intelligence. The goal of this book is to provide C# programmers with practical guidance in solving complex computational challenges using neural networks and C# libraries such as CNTK, and TensorFlowSharp. This book will take you on a step-by-step practical journey, covering everything from the mathematical and theoretical aspects of neural networks, to building your own deep neural networks into your applications with the C# and .NET frameworks. This book begins by giving you a quick refresher of neural networks. You will learn how to build a neural network from scratch using packages such as Encog, Aforge, and Accord. You will learn about various concepts and techniques, such as deep networks, perceptrons, optimization algorithms, convolutional networks, and autoencoders. You will learn ways to add intelligent features to your .NET apps, such as facial and motion detection, object detection and labeling, language understanding, knowledge, and intelligent search. Throughout this book, you will be working on interesting demonstrations that will make it easier to implement complex neural networks in your enterprise applications. What you will learn Understand perceptrons and how to implement them in C# Learn how to train and visualize a neural network using cognitive services Perform image recognition for detecting and labeling objects using C# and TensorFlowSharp Detect specific image characteristics such as a face using Accord.Net Demonstrate particle swarm optimization using a

simple XOR problem and Encog Train convolutional neural networks using ConvNetSharp Find optimal parameters for your neural network functions using numeric and heuristic optimization techniques. Who this book is for This book is for Machine Learning Engineers, Data Scientists, Deep Learning Aspirants and Data Analysts who are now looking to move into advanced machine learning and deep learning with C#. Prior knowledge of machine learning and working experience with C# programming is required to take most out of this book

Neural Network Programming with Java - Second Edition

Explore various approaches to organize and extract useful text from unstructured data using Java Key Features Use deep learning and NLP techniques in Java to discover hidden insights in text Work with popular Java libraries such as CoreNLP, OpenNLP, and Mallet Explore machine translation, identifying parts of speech, and topic modeling Book Description Natural Language Processing (NLP) allows you to take any sentence and identify patterns, special names, company names, and more. The second edition of Natural Language Processing with Java teaches you how to perform language analysis with the help of Java libraries, while constantly gaining insights from the outcomes. You'll start by understanding how NLP and its various concepts work. Having got to grips with the basics, you'll explore important tools and libraries in Java for NLP, such as CoreNLP, OpenNLP, Neuroph, and Mallet. You'll then start performing NLP on different inputs and tasks, such as tokenization, model training, parts-of-speech and parsing trees. You'll learn about statistical machine translation, summarization, dialog systems, complex searches, supervised and unsupervised NLP, and more. By the end of this book, you'll have learned more about NLP, neural networks, and various other trained models in Java for enhancing the performance of NLP applications. What you will learn Understand basic NLP tasks and how they relate to one another Discover and use the available tokenization engines Apply search techniques to find people, as well as things, within a document Construct solutions to identify parts of speech within sentences Use parsers to extract relationships between elements of a document Identify topics in a set of documents Explore topic modeling from a document Who this book is for Natural Language Processing with Java is for you if you are a data analyst, data scientist, or machine learning engineer who wants to extract information from a language using Java. Knowledge of Java programming is needed, while a basic understanding of statistics will be useful but not mandatory.

[Copyright: 3e3b2b53210a0221538837d84cf41e40](#)