Manning Deep Learning With Python

This is likewise one of the factors by obtaining the soft documents of this manning deep learning with python by online. You might not require more epoch to spend to go to the ebook introduction as skillfully as search for them. In some cases, you likewise complete not discover the statement manning deep learning with python that you are looking for. It will definitely squander the time.

However below, like you visit this web page, it will be thus agreed easy to get as competently as download guide manning deep learning with python

It will not say yes many epoch as we accustom before. You can realize it even though acquit yourself something else at home and even in your workplace. fittingly easy! So, are you question? Just exercise just what we find the money for below as competently as evaluation manning deep learning with python what you bearing in mind to read!

Deep Learning with Python (Book Review) Deep Learning with Python, TensorFlow, and Keras tutorial Favorite AI/ML Books: Intro to ML with Python (Book Review) Learn TensorFlow and Deep Learning fundamentals with Python (code-first introduction) Part 1/2 Is this still the best book on Machine Learning? Building a neural network FROM SCRATCH (no Tensorflow/Pytorch, just numpy \u0026 math) TensorFlow 2.0 Complete Course - Python Neural Networks for Beginners Tutorial Reinforcement Learning in 3 Hours | Full Course using Python Theory of Neural Networks - Deep Learning Without Frameworks These books will help you learn machine learning

Deep Learning With Python | Deep Learning Tutorial For Beginners | EdurekaWhy You Shouldn't Learn Python In 2021 Best Books For Python Why You Should NOT Learn Machine Learning! I'm just not that good at coding You should put this in all your Python scripts Python Tutorial - 12.1 - Install Python Module (using pip) PCEP-30-01 | Exam Passed - Here are my thoughts. | Certified Entry-Level Python Programmer A Trump Speech Written By Artificial Intelligence | The New Yorker

installation of library in anaconda, python , pip, conda and navigatorPython Machine Learning Tutorial #1 - Introduction Best Machine Learning Books Python Machine Learning Tutorial (Data Science) Machine Learning Projects You NEVER Know Existed Best Books for Neural Networks or Deep Learning

Neural Networks from Scratch - P.1 Intro and Neuron CodeBut what is a neural network? | Chapter 1, Deep learning PACKT VS O'REILLY. Which learning platform is better? You'll be SURPRISED by the answer!

Have you read these FANTASTIC PYTHON BOOKS? LEARN PYTHON!Manning Deep Learning With Python
In the podcast, Rosaria Silipo talks about the emerging trends in deep learning, with focus on low code visual programming to help data scientists apply deep learning techniques without having to ...

Big Data Processing with Apache Spark - Part 4: Spark Machine Learning
We refer the reader to Schmidhuber 15 for a comprehensive review of these innovations. These methods form the foundation of deep learning, a subfield of ML built on neural networks. Neural networks ...

Machine Learning in Oncology: Methods, Applications, and Challenges
and 1 set of problem solving tasks using Python (which accounts for 30% of the final assessment). The course will cover several theoretical concepts, such as optimisation methods for learning deep ...

Artificial Intelligence
Introduction to machine learning and artificial neural networks; feasibility of learning; deep feedforward networks; regularization in deep learning; optimization of deep learning models; ...

ENSC 813/413 - Deep Learning Systems in Engineering
This new release adds cross-architecture software development capabilities for CPUs and GPUs through the first unified C++/SYCL/Fortran compiler and Data Parallel Python and expands ... to facilitate ...

Intel Innovation Spotlights New Products, Technology and Tools for Developers
Other indicators that are not linked to their educational and language skills are if they have difficulty getting dressed properly, if they enjoy being read to but show no interest in learning ...

How to spot dyslexia in your child
In this class we will apply methods such as generalized linear models, multi-level models, unsupervised learning, and basic neural networks to biological problems. Hands-on activities using Python ...

Course Listing in Biological Sciences
Victoria's great aunt, Marie-T herese Kouao, and her boyfriend, Carl Manning, w ere both convicted of her m u rder in February 2000 and jailed for life.

Social workers in tragic Climbie case to get £4,000 rise
The concepts discussed in the lectures are analyzed in deep and applied through practical projects ... fortran, Java, BASIC, or Python. An overview of all aspects of wind energy power generation: The ...

Course Listing for Mechanical Engineering
This advertisement has not loaded yet, but your article continues below. Banksy supplied the art on one condition — 'it would be his hero Christopher holding the roller' 'This was my first ...

Entertainment News | Breaking Headlines, Gossip, Photos & Videos | National Post
We refer the reader to Schmidhuber 15 for a comprehensive review of these innovations. These methods form the foundation of deep learning, a subfield of ML built on neural networks. Neural networks ...

Machine Learning in Oncology: Methods, Applications, and Challenges
SAN FRANCISCO, October 27, 2021--At Intel Innovation, the company returns to its developer roots, highlighting a renewed commitment to the community and a developer-first approach.

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

"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

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

Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interestingor useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1 CLASSIFICATION Machine learning basics Classifying with k-Nearest Neighbors Splitting datasets one feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines Improving classification with the AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify data Simplifying data with the singular value decomposition Big data and MapReduce

Summary Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. The book builds your understanding of deep learning through intuitive explanations and practical examples. Continue your journey into the world of deep learning with Deep Learning with R in Motion, a practical, hands-on video course available exclusively at Manning.com (www.manning.com/livevideo/deep-learning-with-r-in-motion). Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. Deep-learning systems now enable previously impossible smart applications, revolutionizing image recognition and natural-language processing, and identifying complex patterns in data. The Keras deep-learning library provides data scientists and developers working in R a state-of-the-art toolset for tackling deep-learning tasks. About the Book Deep Learning with R introduces the world of deep learning using the powerful Keras library and its R language interface. Initially written for Python by Keras creator and Google AI researcher François Chollet and adapted for R by RStudio founder J. J. Allaire, this book builds your understanding of deep learning through intuitive explanations and practical examples. You'll practice your new skills with R-based applications in computer vision, natural-language processing, and generative models. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image classification and generation Deep learning for text and sequences About the Reader You'll need intermediate R programming skills. No previous experience with machine learning or deep learning is assumed. About the Authors François Chollet is a deep-learning researcher at Google and the author of the Keras library. J.J. Allaire is the founder of RStudio and the author of the R interfaces to TensorFlow and Keras. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions

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

Probabilistic Deep Learning is a hands-on guide to the principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications. Summary Probabilistic Deep Learning: With Python, Keras and TensorFlow Probability teaches the increasingly popular probabilistic approach to deep learning that allows you to refine your results more quickly and accurately without much trial-and-error testing. Emphasizing practical techniques that use the Python-based Tensorflow Probability Framework, you'll learn to build highly-performant deep learning applications that can reliably handle the noise and uncertainty of real-world data. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology The world is a noisy and uncertain place. Probabilistic deep learning models capture that noise and uncertainty, pulling it into real-world scenarios. Crucial for self-driving cars and scientific testing, these techniques help deep learning engineers assess the accuracy of their results, spot errors, and improve their understanding of how algorithms work. About the book Probabilistic Deep Learning is a hands-on guide to the principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications. What's inside Explore maximum likelihood and the statistical basis of deep learning Discover probabilistic models that can indicate possible outcomes Learn to use normalizing flows for modeling and generating complex distributions Use Bayesian neural networks to access the uncertainty in the model About the reader For experienced machine learning developers. About the author Oliver Dürr is a professor at the University of Applied Sciences in Konstanz, Germany. Beate Sick holds a chair for applied statistics at ZHAW and works as a researcher and lecturer at the University of Zurich. Elvis Murina is a data scientist. Table of Contents PART 1 - BASICS OF DEEP LEARNING 1 Introduction to probabilistic deep learning 2 Neural network architectures 3 Principles of curve fitting PART 2 - MAXIMUM LIKELIHOOD APPROACHES FOR PROBABILISTIC DL MODELS 4 Building loss functions with the likelihood approach 5 Probabilistic deep learning models with TensorFlow Probability 6 Probabilistic deep learning models in the wild PART 3 - BAYESIAN APPROACHES FOR PROBABILISTIC DL MODELS 7 Bayesian learning 8 Bayesian neural networks

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 Shanqing 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

Time to flex your machine learning muscles! Take on the carefully designed challenges of the Machine Learning Bookcamp and master essential ML techniques through practical application. Summary In Machine Learning Bookcamp you will: Collect and clean data for training models Use popular Python tools, including NumPy, Scikit-Learn, and TensorFlow Apply ML to complex datasets with images Deploy ML models to a production-ready environment The only way to learn is to practice! In Machine Learning Bookcamp, you'll create and deploy Python-based machine learning models for a variety of increasingly challenging projects. Taking you from the basics of machine learning to complex applications such as image analysis, each new project builds on what you've learned in previous chapters. You'll build a portfolio of business-relevant machine learning projects that hiring managers will be excited to see. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Master key machine learning concepts as you build actual projects! Machine learning is what you need for analyzing customer behavior, predicting price trends, evaluating risk, and much more. To master ML, you need great examples, clear explanations, and lots of practice. This book delivers all three! About the book Machine Learning Bookcamp presents realistic, practical machine learning scenarios, along with crystal-clear coverage of key concepts. In it, you'll complete engaging projects, such as creating a car price predictor using linear regression and deploying a churn prediction service. You'll go beyond the algorithms and explore important techniques like deploying ML applications on serverless systems and serving models with Kubernetes and Kubeflow. Dig in, get your hands dirty, and have fun building your ML skills! What's inside Collect and clean data for training models Use popular Python tools, including NumPy, Scikit-Learn, and TensorFlow Deploy ML models to a production-ready environment About the reader Python programming skills assumed. No previous machine learning knowledge is required. About the author Alexey Grigorev is a principal data scientist at OLX Group. He runs DataTalks.Club, a community of people who love data. Table of Contents 1 Introduction to machine

learning 2 Machine learning for regression 3 Machine learning for classification 4 Evaluation metrics for classification 5 Deploying machine learning models 6 Decision trees and ensemble learning 7 Neural networks and deep learning 8 Serverless deep learning 9 Serving models with Kubernetes and Kubeflow

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