

Deep Learning For Natural Language Processing Sol

This is likewise one of the factors by obtaining the soft documents of this **Deep Learning For Natural Language Processing Sol** by online. You might not require more grow old to spend to go to the books instigation as capably as search for them. In some cases, you likewise realize not discover the proclamation Deep Learning For Natural Language Processing Sol that you are looking for. It will unquestionably squander the time.

However below, once you visit this web page, it will be therefore very easy to get as well as download guide Deep Learning For Natural Language Processing Sol

It will not assume many get older as we accustom before. You can do it though affect something else at house and even in your workplace. correspondingly easy! So, are you question? Just exercise just what we have enough money below as competently as review **Deep Learning For Natural Language Processing Sol** what you like to read!

Practical Natural Language Processing - Sowmya Vajjala 2020-06-17

Many books and courses tackle natural language processing (NLP) problems with toy use cases and well-defined datasets. But if you want to build, iterate, and scale NLP systems in a business setting and tailor them for particular industry verticals, this is your guide. Software engineers and data scientists will learn how to navigate the maze of options available at each step of the journey. Through the course of the book, authors Sowmya Vajjala, Bodhisattwa Majumder, Anuj Gupta, and Harshit Surana will guide you through the process of building real-world NLP solutions embedded in larger product setups. You'll learn how to adapt your solutions for different industry verticals such as healthcare, social media, and retail. With this book, you'll: Understand the wide spectrum of problem statements, tasks, and solution approaches within NLP Implement and evaluate different NLP applications using machine learning and deep learning methods Fine-tune your NLP solution based on your business problem and industry vertical Evaluate various algorithms and approaches for NLP product tasks, datasets, and stages Produce software solutions following best practices around release, deployment, and DevOps for NLP systems Understand best practices, opportunities, and the roadmap for NLP from a business and product leader's perspective

IBM Watson Solutions for Machine Learning - Arindam Ganguly 2021-06-19

Utilize Python and IBM Watson to put real-life use cases into production. KEY FEATURES ● Use of popular Python packages for building Machine Learning solutions from scratch. ● Practice various IBM Watson Machine Learning tools for Computer Vision and Natural Language Processing applications. ● Expert-led best practices to put your Machine Learning solutions into the production environment. DESCRIPTION This book will take you through the journey of some amazing tools IBM Watson has to offer to leverage your machine learning concepts to solve some real-life use cases that are pertinent to the current industry. This book explores the various Machine Learning fundamental concepts and how to use the Python programming language to deal with real-world use cases. It explains how to take your code and deploy it into IBM Cloud leveraging IBM Watson Machine Learning. While doing so, the book also introduces you to several amazing IBM Watson tools such as Watson Assistant, Watson Discovery, and Watson Visual Recognition to ease out various machine learning tasks such as building a chatbot, creating a natural language processing pipeline, or an optical object detection application without a single line of code. It covers Watson Auto AI with which you can apply various machine learning algorithms and pick out the best for your dataset without a single line of code. Finally, you will be able to deploy all of these into IBM Cloud and configure your application to maintain the production-level runtime. After reading this book, you will find yourself confident to administer any machine learning use case and deploy it into production without any hassle. You will be able to take up a complete end-to-end machine learning project with complete responsibility and deliver the best standards the current industry has to offer. Towards the end of this book, you will be able to build an end-to-end production-level application and deploy it into Cloud. WHAT YOU WILL LEARN ● Review the basics of Machine Learning and learn implementation using Python. ● Learn deployment using IBM Watson Studio and Watson Machine Learning. ● Learn how to use Watson Auto AI to automate hyperparameter tuning. ● Learn Watson Assistant, Watson Visual Recognition, and Watson Discovery. ● Learn how to implement the various layers of an end-to-end AI application. ● Learn all the

configurations needed for production deployment to Cloud. WHO THIS BOOK IS FOR This book is for all data professionals, ML enthusiasts, and software developers who are looking for real solutions to be developed. The reader is expected to have a prior knowledge of the web application architecture and basic Python fundamentals. TABLE OF CONTENTS 1. Introduction to Machine Learning 2. Deep Learning 3. Features and Metrics 4. Build Your Own Chatbot 5. First Complete Machine Learning Project 6. Perfecting Our Model 7. Visual Recognition 8. Watson Discovery 9. Deployment and Others 10. Deploying the Food Ordering Bot

Introduction to Natural Language Processing - Jacob Eisenstein 2019-10-01

A survey of computational methods for understanding, generating, and manipulating human language, which offers a synthesis of classical representations and algorithms with contemporary machine learning techniques. This textbook provides a technical perspective on natural language processing—methods for building computer software that understands, generates, and manipulates human language. It emphasizes contemporary data-driven approaches, focusing on techniques from supervised and unsupervised machine learning. The first section establishes a foundation in machine learning by building a set of tools that will be used throughout the book and applying them to word-based textual analysis. The second section introduces structured representations of language, including sequences, trees, and graphs. The third section explores different approaches to the representation and analysis of linguistic meaning, ranging from formal logic to neural word embeddings. The final section offers chapter-length treatments of three transformative applications of natural language processing: information extraction, machine translation, and text generation. End-of-chapter exercises include both paper-and-pencil analysis and software implementation. The text synthesizes and distills a broad and diverse research literature, linking contemporary machine learning techniques with the field's linguistic and computational foundations. It is suitable for use in advanced undergraduate and graduate-level courses and as a reference for software engineers and data scientists. Readers should have a background in computer programming and college-level mathematics. After mastering the material presented, students will have the technical skill to build and analyze novel natural language processing systems and to understand the latest research in the field.

Deep Learning in Natural Language Processing - Li Deng 2018-05-23

In recent years, deep learning has fundamentally changed the landscapes of a number of areas in artificial intelligence, including speech, vision, natural language, robotics, and game playing. In particular, the striking success of deep learning in a wide variety of natural language processing (NLP) applications has served as a benchmark for the advances in one of the most important tasks in artificial intelligence. This book reviews the state of the art of deep learning research and its successful applications to major NLP tasks, including speech recognition and understanding, dialogue systems, lexical analysis, parsing, knowledge graphs, machine translation, question answering, sentiment analysis, social computing, and natural language generation from images. Outlining and analyzing various research frontiers of NLP in the deep learning era, it features self-contained, comprehensive chapters written by leading researchers in the field. A glossary of technical terms and commonly used acronyms in the intersection of deep learning and NLP is also provided. The book appeals to advanced undergraduate and graduate students, post-doctoral

researchers, lecturers and industrial researchers, as well as anyone interested in deep learning and natural language processing.

Natural Language Processing with Python - Steven Bird 2009-06-12

This book offers a highly accessible introduction to natural language processing, the field that supports a variety of language technologies, from predictive text and email filtering to automatic summarization and translation. With it, you'll learn how to write Python programs that work with large collections of unstructured text. You'll access richly annotated datasets using a comprehensive range of linguistic data structures, and you'll understand the main algorithms for analyzing the content and structure of written communication. Packed with examples and exercises, *Natural Language Processing with Python* will help you: Extract information from unstructured text, either to guess the topic or identify "named entities" Analyze linguistic structure in text, including parsing and semantic analysis Access popular linguistic databases, including WordNet and treebanks Integrate techniques drawn from fields as diverse as linguistics and artificial intelligence This book will help you gain practical skills in natural language processing using the Python programming language and the Natural Language Toolkit (NLTK) open source library. If you're interested in developing web applications, analyzing multilingual news sources, or documenting endangered languages -- or if you're simply curious to have a programmer's perspective on how human language works -- you'll find *Natural Language Processing with Python* both fascinating and immensely useful.

Blueprints for Text Analytics Using Python - Jens Albrecht 2020-12-04

Turning text into valuable information is essential for businesses looking to gain a competitive advantage. With recent improvements in natural language processing (NLP), users now have many options for solving complex challenges. But it's not always clear which NLP tools or libraries would work for a business's needs, or which techniques you should use and in what order. This practical book provides data scientists and developers with blueprints for best practice solutions to common tasks in text analytics and natural language processing. Authors Jens Albrecht, Sidharth Ramachandran, and Christian Winkler provide real-world case studies and detailed code examples in Python to help you get started quickly. Extract data from APIs and web pages Prepare textual data for statistical analysis and machine learning Use machine learning for classification, topic modeling, and summarization Explain AI models and classification results Explore and visualize semantic similarities with word embeddings Identify customer sentiment in product reviews Create a knowledge graph based on named entities and their relations

[Deep Learning for Natural Language Processing](#) - Mihai Surdeanu 2022-06

[Deep Learning and Edge Computing Solutions for High Performance Computing](#) - A. Suresh 2021-01-27

This book provides an insight into ways of inculcating the need for applying mobile edge data analytics in bioinformatics and medicine. The book is a comprehensive reference that provides an overview of the current state of medical treatments and systems and offers emerging solutions for a more personalized approach to the healthcare field. Topics include deep learning methods for applications in object detection and identification, object tracking, human action recognition, and cross-modal and multimodal data analysis. High performance computing systems for applications in healthcare are also discussed. The contributors also include information on microarray data analysis, sequence analysis, genomics based analytics, disease network analysis, and techniques for big data Analytics and health information technology.

Deep Learning Research Applications for Natural Language Processing - Ashok Kumar, L. 2022-12-09

Humans have the most advanced method of communication, which is known as natural language. While humans can use computers to send voice and text messages to each other, computers do not innately know how to process natural language. In recent years, deep learning has primarily transformed the perspectives of a variety of fields in artificial intelligence (AI), including speech, vision, and natural language processing (NLP). The extensive success of deep learning in a wide variety of applications has served as a benchmark for the many downstream tasks in AI. The field of computer vision has taken great leaps in recent years and surpassed humans in tasks related to detecting and labeling objects thanks to advances in deep learning and neural networks. *Deep Learning Research Applications for Natural Language Processing* explains the

concepts and state-of-the-art research in the fields of NLP, speech, and computer vision. It provides insights into using the tools and libraries in Python for real-world applications. Covering topics such as deep learning algorithms, neural networks, and advanced prediction, this premier reference source is an excellent resource for computational linguists, software engineers, IT managers, computer scientists, students and faculty of higher education, libraries, researchers, and academicians.

Healthcare Solutions Using Machine Learning and Informatics - Punit Gupta 2022-10-28

Healthcare Solutions Using Machine Learning and Informatics covers novel and innovative solutions for healthcare that apply machine learning and biomedical informatics technology. The healthcare sector is one of the most critical in society. This book presents a series of artificial intelligence, machine learning, and intelligent IoT-based solutions for medical image analysis, medical big-data processing, and disease predictions. Machine learning and artificial intelligence use cases in healthcare presented in the book give researchers, practitioners, and students a wide range of practical examples of cross-domain convergence. The wide variety of topics covered include: Artificial Intelligence in healthcare Machine learning solutions for such disease as diabetes, arthritis, cardiovascular disease, and COVID-19 Big data analytics solutions for healthcare data processing Reliable biomedical applications using AI models Intelligent IoT in healthcare The book explains fundamental concepts as well as the advanced use cases, illustrating how to apply emerging technologies such as machine learning, AI models, and data informatics into practice to tackle challenges in the field of healthcare with real-world scenarios. Chapters contributed by noted academicians and professionals examine various solutions, frameworks, applications, case studies, and best practices in the healthcare domain.

[Deep Learning for NLP and Speech Recognition](#) - Uday Kamath 2019-06-10

This textbook explains Deep Learning Architecture, with applications to various NLP Tasks, including Document Classification, Machine Translation, Language Modeling, and Speech Recognition. With the widespread adoption of deep learning, natural language processing (NLP), and speech applications in many areas (including Finance, Healthcare, and Government) there is a growing need for one comprehensive resource that maps deep learning techniques to NLP and speech and provides insights into using the tools and libraries for real-world applications. *Deep Learning for NLP and Speech Recognition* explains recent deep learning methods applicable to NLP and speech, provides state-of-the-art approaches, and offers real-world case studies with code to provide hands-on experience. Many books focus on deep learning theory or deep learning for NLP-specific tasks while others are cookbooks for tools and libraries, but the constant flux of new algorithms, tools, frameworks, and libraries in a rapidly evolving landscape means that there are few available texts that offer the material in this book. The book is organized into three parts, aligning to different groups of readers and their expertise. The three parts are: Machine Learning, NLP, and Speech Introduction The first part has three chapters that introduce readers to the fields of NLP, speech recognition, deep learning and machine learning with basic theory and hands-on case studies using Python-based tools and libraries. Deep Learning Basics The five chapters in the second part introduce deep learning and various topics that are crucial for speech and text processing, including word embeddings, convolutional neural networks, recurrent neural networks and speech recognition basics. Theory, practical tips, state-of-the-art methods, experimentations and analysis in using the methods discussed in theory on real-world tasks. Advanced Deep Learning Techniques for Text and Speech The third part has five chapters that discuss the latest and cutting-edge research in the areas of deep learning that intersect with NLP and speech. Topics including attention mechanisms, memory augmented networks, transfer learning, multi-task learning, domain adaptation, reinforcement learning, and end-to-end deep learning for speech recognition are covered using case studies.

[Practical Natural Language Processing](#) - Sowmya Vajjala 2020-06-17

Many books and courses tackle natural language processing (NLP) problems with toy use cases and well-defined datasets. But if you want to build, iterate, and scale NLP systems in a business setting and tailor them for particular industry verticals, this is your guide. Software engineers and data scientists will learn how to navigate the maze of options available at each step of the journey. Through the course of the book, authors Sowmya Vajjala, Bodhisattwa Majumder, Anuj Gupta, and Harshit Surana will guide you through the process of building real-world NLP solutions embedded in larger product setups. You'll learn how to

adapt your solutions for different industry verticals such as healthcare, social media, and retail. With this book, you'll: Understand the wide spectrum of problem statements, tasks, and solution approaches within NLP Implement and evaluate different NLP applications using machine learning and deep learning methods Fine-tune your NLP solution based on your business problem and industry vertical Evaluate various algorithms and approaches for NLP product tasks, datasets, and stages Produce software solutions following best practices around release, deployment, and DevOps for NLP systems Understand best practices, opportunities, and the roadmap for NLP from a business and product leader's perspective [Representation Learning for Natural Language Processing](#) - Zhiyuan Liu 2020-07-03

This open access book provides an overview of the recent advances in representation learning theory, algorithms and applications for natural language processing (NLP). It is divided into three parts. Part I presents the representation learning techniques for multiple language entries, including words, phrases, sentences and documents. Part II then introduces the representation techniques for those objects that are closely related to NLP, including entity-based world knowledge, sememe-based linguistic knowledge, networks, and cross-modal entries. Lastly, Part III provides open resource tools for representation learning techniques, and discusses the remaining challenges and future research directions. The theories and algorithms of representation learning presented can also benefit other related domains such as machine learning, social network analysis, semantic Web, information retrieval, data mining and computational biology. This book is intended for advanced undergraduate and graduate students, post-doctoral fellows, researchers, lecturers, and industrial engineers, as well as anyone interested in representation learning and natural language processing.

Natural Language Processing in Artificial Intelligence - Brojo Kishore Mishra 2020-11-01

This volume focuses on natural language processing, artificial intelligence, and allied areas. Natural language processing enables communication between people and computers and automatic translation to facilitate easy interaction with others around the world. This book discusses theoretical work and advanced applications, approaches, and techniques for computational models of information and how it is presented by language (artificial, human, or natural) in other ways. It looks at intelligent natural language processing and related models of thought, mental states, reasoning, and other cognitive processes. It explores the difficult problems and challenges related to partiality, underspecification, and context-dependency, which are signature features of information in nature and natural languages. Key features:

- Addresses the functional frameworks and workflow that are trending in NLP and AI
- Looks at the latest technologies and the major challenges, issues, and advances in NLP and AI
- Explores an intelligent field monitoring and automated system through AI with NLP and its implications for the real world
- Discusses data acquisition and presents a real-time case study with illustrations related to data-intensive technologies in AI and NLP

Natural Language Processing - Yue Zhang 2021-01-07

This undergraduate textbook introduces essential machine learning concepts in NLP in a unified and gentle mathematical framework.

Data Science for Healthcare - Sergio Consoli 2019-02-23

This book seeks to promote the exploitation of data science in healthcare systems. The focus is on advancing the automated analytical methods used to extract new knowledge from data for healthcare applications. To do so, the book draws on several interrelated disciplines, including machine learning, big data analytics, statistics, pattern recognition, computer vision, and Semantic Web technologies, and focuses on their direct application to healthcare. Building on three tutorial-like chapters on data science in healthcare, the following eleven chapters highlight success stories on the application of data science in healthcare, where data science and artificial intelligence technologies have proven to be very promising. This book is primarily intended for data scientists involved in the healthcare or medical sector. By reading this book, they will gain essential insights into the modern data science technologies needed to advance innovation for both healthcare businesses and patients. A basic grasp of data science is recommended in order to fully benefit from this book.

Cognitive Computing Recipes - Adnan Masood 2019-03-27

Solve your AI and machine learning problems using complete and real-world code examples. Using a problem-solution approach, this book makes deep learning and machine learning accessible to everyday

developers, by providing a combination of tools such as cognitive services APIs, machine learning platforms, and libraries. Along with an overview of the contemporary technology landscape, Machine Learning and Deep Learning with Cognitive Computing Recipes covers the business case for machine learning and deep learning. Covering topics such as digital assistants, computer vision, text analytics, speech, and robotics process automation this book offers a comprehensive toolkit that you can apply quickly and easily in your own projects. With its focus on Microsoft Cognitive Services offerings, you'll see recipes using multiple different environments including TensorFlow and CNTK to give you a broader perspective of the deep learning ecosystem. What You Will Learn Build production-ready solutions using Microsoft Cognitive Services APIs Apply deep learning using TensorFlow and Microsoft Cognitive Toolkit (CNTK) Solve enterprise problems in natural language processing and computer vision Discover the machine learning development life cycle - from formal problem definition to deployment at scale Who This Book Is For Software engineers and enterprise architects who wish to understand machine learning and deep learning by building applications and solving real-world business problems.

Hands-On Natural Language Processing with Python - Rajesh Arumugam 2018-07-18

Foster your NLP applications with the help of deep learning, NLTK, and TensorFlow Key Features Weave neural networks into linguistic applications across various platforms Perform NLP tasks and train its models using NLTK and TensorFlow Boost your NLP models with strong deep learning architectures such as CNNs and RNNs Book Description Natural language processing (NLP) has found its application in various domains, such as web search, advertisements, and customer services, and with the help of deep learning, we can enhance its performances in these areas. Hands-On Natural Language Processing with Python teaches you how to leverage deep learning models for performing various NLP tasks, along with best practices in dealing with today's NLP challenges. To begin with, you will understand the core concepts of NLP and deep learning, such as Convolutional Neural Networks (CNNs), recurrent neural networks (RNNs), semantic embedding, Word2vec, and more. You will learn how to perform each and every task of NLP using neural networks, in which you will train and deploy neural networks in your NLP applications. You will get accustomed to using RNNs and CNNs in various application areas, such as text classification and sequence labeling, which are essential in the application of sentiment analysis, customer service chatbots, and anomaly detection. You will be equipped with practical knowledge in order to implement deep learning in your linguistic applications using Python's popular deep learning library, TensorFlow. By the end of this book, you will be well versed in building deep learning-backed NLP applications, along with overcoming NLP challenges with best practices developed by domain experts. What you will learn Implement semantic embedding of words to classify and find entities Convert words to vectors by training in order to perform arithmetic operations Train a deep learning model to detect classification of tweets and news Implement a question-answer model with search and RNN models Train models for various text classification datasets using CNN Implement WaveNet a deep generative model for producing a natural-sounding voice Convert voice-to-text and text-to-voice Train a model to convert speech-to-text using DeepSpeech Who this book is for Hands-on Natural Language Processing with Python is for you if you are a developer, machine learning or an NLP engineer who wants to build a deep learning application that leverages NLP techniques. This comprehensive guide is also useful for deep learning users who want to extend their deep learning skills in building NLP applications. All you need is the basics of machine learning and Python to enjoy the book.

Natural Language Processing Crash Course for Beginners - Ai Publishing 2020-08-04

Natural Language Processing Crash Course for Beginners Artificial Intelligence (AI) isn't the latest fad! The reason is AI has been around since 1956, and its relevance is evident in every field today. Artificial Intelligence incorporates human intelligence into machines. Machine Learning (ML), a branch of AI, enables machines to learn by themselves. Deep Learning (DL), a subfield of Machine Learning, uses algorithms that are inspired by the functioning of the human brain. Natural Language Processing (NLP) combines computational linguistics and Artificial Intelligence, enabling computers and humans to communicate seamlessly. And NLP is immensely powerful and impactful as every business is looking to integrate it into their day to day dealings. How Is This Book Different? This book by AI Publishing is carefully crafted, giving equal importance to the theoretical concepts as well as the practical aspects of

natural language processing. In each chapter of the second half of the book, the theoretical concepts of different types of deep learning and NLP techniques have been covered in-depth, followed by practical examples. You will learn how to apply different NLP techniques using the TensorFlow and Keras libraries for Python. Each chapter contains exercises that are designed to evaluate your understanding of the concepts covered in that chapter. Also, in the Resources section of each chapter, you can access the Python notebook. The author has also compiled a list of hands-on NLP projects and competitions that you can try on your own. The main benefit of purchasing this book is you get immediate access to all the extra learning material presented with this book--Python codes, exercises, PDFs, and references--on the publisher's website without having to spend an extra cent. You can download the datasets used in this book at runtime, or you can access them in the Resources/Datasets folder. The author holds your hand through everything. He provides you a step by step explanation of the installation of the software needed to implement the various NLP techniques in this book. You can start experimenting with the practical aspects of NLP right from the beginning. Even if you are new to Python, you'll find the ultra-short course on Python programming language in the second chapter immensely helpful. You get all the codes and datasets with this book. So, if you have access to a computer with the internet, you can get started. The topics covered include: What is Natural Language Processing? Environment Setup and Python Crash Course Introduction to Deep Learning Text Cleaning and Manipulation Common NLP Tasks Importing Text Data from Various Sources Word Embeddings: Converting Words to Numbers IMDB Movies Sentimental Analysis Ham and Spam Message Classification Text Summarization and Topic Modeling Text Classification with Deep Learning Text Translation Using Seq2Seq Model State of the Art NLP with BERT Transformers Hands-on NLP Projects/Articles for Practice Exercise Solutions Click the BUY button and download the book now to start your Natural Language Processing journey.

Python Natural Language Processing - Jalaj Thanaki 2017-07-31

Leverage the power of machine learning and deep learning to extract information from text data About This Book Implement Machine Learning and Deep Learning techniques for efficient natural language processing Get started with NLTK and implement NLP in your applications with ease Understand and interpret human languages with the power of text analysis via Python Who This Book Is For This book is intended for Python developers who wish to start with natural language processing and want to make their applications smarter by implementing NLP in them. What You Will Learn Focus on Python programming paradigms, which are used to develop NLP applications Understand corpus analysis and different types of data attribute. Learn NLP using Python libraries such as NLTK, Polyglot, SpaCy, Stanford CoreNLP and so on Learn about Features Extraction and Feature selection as part of Features Engineering. Explore the advantages of vectorization in Deep Learning. Get a better understanding of the architecture of a rule-based system. Optimize and fine-tune Supervised and Unsupervised Machine Learning algorithms for NLP problems. Identify Deep Learning techniques for Natural Language Processing and Natural Language Generation problems. In Detail This book starts off by laying the foundation for Natural Language Processing and why Python is one of the best options to build an NLP-based expert system with advantages such as Community support, availability of frameworks and so on. Later it gives you a better understanding of available free forms of corpus and different types of dataset. After this, you will know how to choose a dataset for natural language processing applications and find the right NLP techniques to process sentences in datasets and understand their structure. You will also learn how to tokenize different parts of sentences and ways to analyze them. During the course of the book, you will explore the semantic as well as syntactic analysis of text. You will understand how to solve various ambiguities in processing human language and will come across various scenarios while performing text analysis. You will learn the very basics of getting the environment ready for natural language processing, move on to the initial setup, and then quickly understand sentences and language parts. You will learn the power of Machine Learning and Deep Learning to extract information from text data. By the end of the book, you will have a clear understanding of natural language processing and will have worked on multiple examples that implement NLP in the real world. Style and approach This book teaches the readers various aspects of natural language Processing using NLTK. It takes the reader from the basic to advance level in a smooth way.

Deep Learning for Natural Language Processing - Palash Goyal 2018-06-26

Discover the concepts of deep learning used for natural language processing (NLP), with full-fledged examples of neural network models such as recurrent neural networks, long short-term memory networks, and sequence-2-sequence models. You'll start by covering the mathematical prerequisites and the fundamentals of deep learning and NLP with practical examples. The first three chapters of the book cover the basics of NLP, starting with word-vector representation before moving onto advanced algorithms. The final chapters focus entirely on implementation, and deal with sophisticated architectures such as RNN, LSTM, and Seq2seq, using Python tools: TensorFlow, and Keras. Deep Learning for Natural Language Processing follows a progressive approach and combines all the knowledge you have gained to build a question-answer chatbot system. This book is a good starting point for people who want to get started in deep learning for NLP. All the code presented in the book will be available in the form of IPython notebooks and scripts, which allow you to try out the examples and extend them in interesting ways. What You Will Learn Gain the fundamentals of deep learning and its mathematical prerequisites Discover deep learning frameworks in Python Develop a chatbot Implement a research paper on sentiment classification Who This Book Is For Software developers who are curious to try out deep learning with NLP.

Machine Learning and Knowledge Extraction - Andreas Holzinger 2019-08-22

This book constitutes the refereed proceedings of the IFIP TC 5, TC 12, WG 8.4, 8.9, 12.9 International Cross-Domain Conference for Machine Learning and Knowledge Extraction, CD-MAKE 2019, held in Canterbury, UK, in August 2019. The 25 revised full papers presented were carefully reviewed and selected from 45 submissions. The cross-domain integration and appraisal of different fields provides an atmosphere to foster different perspectives and opinions; it will offer a platform for novel ideas and a fresh look on the methodologies to put these ideas into business for the benefit of humanity.

Deep Natural Language Processing and AI Applications for Industry 5.0 - Tanwar, Poonam 2021-06-25

To sustain and stay at the top of the market and give absolute comfort to the consumers, industries are using different strategies and technologies. Natural language processing (NLP) is a technology widely penetrating the market, irrespective of the industry and domains. It is extensively applied in businesses today, and it is the buzzword in every engineer's life. NLP can be implemented in all those areas where artificial intelligence is applicable either by simplifying the communication process or by refining and analyzing information. Neural machine translation has improved the imitation of professional translations over the years. When applied in neural machine translation, NLP helps educate neural machine networks. This can be used by industries to translate low-impact content including emails, regulatory texts, etc. Such machine translation tools speed up communication with partners while enriching other business interactions. Deep Natural Language Processing and AI Applications for Industry 5.0 provides innovative research on the latest findings, ideas, and applications in fields of interest that fall under the scope of NLP including computational linguistics, deep NLP, web analysis, sentiments analysis for business, and industry perspective. This book covers a wide range of topics such as deep learning, deepfakes, text mining, blockchain technology, and more, making it a crucial text for anyone interested in NLP and artificial intelligence, including academicians, researchers, professionals, industry experts, business analysts, data scientists, data analysts, healthcare system designers, intelligent system designers, practitioners, and students.

Python Deep Learning Cookbook - Indra den Bakker 2017-10-27

Solve different problems in modelling deep neural networks using Python, Tensorflow, and Keras with this practical guide About This Book Practical recipes on training different neural network models and tuning them for optimal performance Use Python frameworks like TensorFlow, Caffe, Keras, Theano for Natural Language Processing, Computer Vision, and more A hands-on guide covering the common as well as the not so common problems in deep learning using Python Who This Book Is For This book is intended for machine learning professionals who are looking to use deep learning algorithms to create real-world applications using Python. Thorough understanding of the machine learning concepts and Python libraries such as NumPy, SciPy and scikit-learn is expected. Additionally, basic knowledge in linear algebra and calculus is desired. What You Will Learn Implement different neural network models in Python Select the best Python framework for deep learning such as PyTorch, Tensorflow, MXNet and Keras Apply tips and tricks related to neural networks internals, to boost learning performances Consolidate machine learning principles and

apply them in the deep learning field Reuse and adapt Python code snippets to everyday problems Evaluate the cost/benefits and performance implication of each discussed solution In Detail Deep Learning is revolutionizing a wide range of industries. For many applications, deep learning has proven to outperform humans by making faster and more accurate predictions. This book provides a top-down and bottom-up approach to demonstrate deep learning solutions to real-world problems in different areas. These applications include Computer Vision, Natural Language Processing, Time Series, and Robotics. The Python Deep Learning Cookbook presents technical solutions to the issues presented, along with a detailed explanation of the solutions. Furthermore, a discussion on corresponding pros and cons of implementing the proposed solution using one of the popular frameworks like TensorFlow, PyTorch, Keras and CNTK is provided. The book includes recipes that are related to the basic concepts of neural networks. All techniques s, as well as classical networks topologies. The main purpose of this book is to provide Python programmers a detailed list of recipes to apply deep learning to common and not-so-common scenarios. Style and approach Unique blend of independent recipes arranged in the most logical manner

Deep Learning Illustrated - Jon Krohn 2019-08-05

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." -Tim Urban, author of Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. Deep Learning Illustrated is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn-with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens-presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Natural Language Processing with AWS AI Services - Mona M 2021-11-26

Work through interesting real-life business use cases to uncover valuable insights from unstructured text using AWS AI services Key FeaturesGet to grips with AWS AI services for NLP and find out how to use them to gain strategic insightsRun Python code to use Amazon Textract and Amazon Comprehend to accelerate business outcomesUnderstand how you can integrate human-in-the-loop for custom NLP use cases with Amazon A2IBook Description Natural language processing (NLP) uses machine learning to extract information from unstructured data. This book will help you to move quickly from business questions to high-performance models in production. To start with, you'll understand the importance of NLP in today's business applications and learn the features of Amazon Comprehend and Amazon Textract to build NLP models using Python and Jupyter Notebooks. The book then shows you how to integrate AI in applications for accelerating business outcomes with just a few lines of code. Throughout the book, you'll cover use cases such as smart text search, setting up compliance and controls when processing confidential documents, real-time text analytics, and much more to understand various NLP scenarios. You'll deploy and monitor scalable NLP models in production for real-time and batch requirements. As you advance, you'll explore strategies for including humans in the loop for different purposes in a document processing

workflow. Moreover, you'll learn best practices for auto-scaling your NLP inference for enterprise traffic. Whether you're new to ML or an experienced practitioner, by the end of this NLP book, you'll have the confidence to use AWS AI services to build powerful NLP applications. What you will learnAutomate various NLP workflows on AWS to accelerate business outcomesUse Amazon Textract for text, tables, and handwriting recognition from images and PDF filesGain insights from unstructured text in the form of sentiment analysis, topic modeling, and more using Amazon ComprehendSet up end-to-end document processing pipelines to understand the role of humans in the loopDevelop NLP-based intelligent search solutions with just a few lines of codeCreate both real-time and batch document processing pipelines using PythonWho this book is for If you're an NLP developer or data scientist looking to get started with AWS AI services to implement various NLP scenarios quickly, this book is for you. It will show you how easy it is to integrate AI in applications with just a few lines of code. A basic understanding of machine learning (ML) concepts is necessary to understand the concepts covered. Experience with Jupyter notebooks and Python will be helpful.

Hands-On Transfer Learning with Python - Dipanjan Sarkar 2018-08-31

Deep learning simplified by taking supervised, unsupervised, and reinforcement learning to the next level using the Python ecosystem Key Features Build deep learning models with transfer learning principles in Python implement transfer learning to solve real-world research problems Perform complex operations such as image captioning neural style transfer Book Description Transfer learning is a machine learning (ML) technique where knowledge gained during training a set of problems can be used to solve other similar problems. The purpose of this book is two-fold; firstly, we focus on detailed coverage of deep learning (DL) and transfer learning, comparing and contrasting the two with easy-to-follow concepts and examples. The second area of focus is real-world examples and research problems using TensorFlow, Keras, and the Python ecosystem with hands-on examples. The book starts with the key essential concepts of ML and DL, followed by depiction and coverage of important DL architectures such as convolutional neural networks (CNNs), deep neural networks (DNNs), recurrent neural networks (RNNs), long short-term memory (LSTM), and capsule networks. Our focus then shifts to transfer learning concepts, such as model freezing, fine-tuning, pre-trained models including VGG, inception, ResNet, and how these systems perform better than DL models with practical examples. In the concluding chapters, we will focus on a multitude of real-world case studies and problems associated with areas such as computer vision, audio analysis and natural language processing (NLP). By the end of this book, you will be able to implement both DL and transfer learning principles in your own systems. What you will learn Set up your own DL environment with graphics processing unit (GPU) and Cloud support Delve into transfer learning principles with ML and DL models Explore various DL architectures, including CNN, LSTM, and capsule networks Learn about data and network representation and loss functions Get to grips with models and strategies in transfer learning Walk through potential challenges in building complex transfer learning models from scratch Explore real-world research problems related to computer vision and audio analysis Understand how transfer learning can be leveraged in NLP Who this book is for Hands-On Transfer Learning with Python is for data scientists, machine learning engineers, analysts and developers with an interest in data and applying state-of-the-art transfer learning methodologies to solve tough real-world problems. Basic proficiency in machine learning and Python is required.

Natural Language Processing with TensorFlow - Thushan Ganegedara 2018-05-31

Write modern natural language processing applications using deep learning algorithms and TensorFlow Key Features Focuses on more efficient natural language processing using TensorFlow Covers NLP as a field in its own right to improve understanding for choosing TensorFlow tools and other deep learning approaches Provides choices for how to process and evaluate large unstructured text datasets Learn to apply the TensorFlow toolbox to specific tasks in the most interesting field in artificial intelligence Book Description Natural language processing (NLP) supplies the majority of data available to deep learning applications, while TensorFlow is the most important deep learning framework currently available. Natural Language Processing with TensorFlow brings TensorFlow and NLP together to give you invaluable tools to work with the immense volume of unstructured data in today's data streams, and apply these tools to specific NLP tasks. Thushan Ganegedara starts by giving you a grounding in NLP and TensorFlow basics.

You'll then learn how to use Word2vec, including advanced extensions, to create word embeddings that turn sequences of words into vectors accessible to deep learning algorithms. Chapters on classical deep learning algorithms, like convolutional neural networks (CNN) and recurrent neural networks (RNN), demonstrate important NLP tasks as sentence classification and language generation. You will learn how to apply high-performance RNN models, like long short-term memory (LSTM) cells, to NLP tasks. You will also explore neural machine translation and implement a neural machine translator. After reading this book, you will gain an understanding of NLP and you'll have the skills to apply TensorFlow in deep learning NLP applications, and how to perform specific NLP tasks. What you will learn Core concepts of NLP and various approaches to natural language processing How to solve NLP tasks by applying TensorFlow functions to create neural networks Strategies to process large amounts of data into word representations that can be used by deep learning applications Techniques for performing sentence classification and language generation using CNNs and RNNs About employing state-of-the-art advanced RNNs, like long short-term memory, to solve complex text generation tasks How to write automatic translation programs and implement an actual neural machine translator from scratch The trends and innovations that are paving the future in NLP Who this book is for This book is for Python developers with a strong interest in deep learning, who want to learn how to leverage TensorFlow to simplify NLP tasks. Fundamental Python skills are assumed, as well as some knowledge of machine learning and undergraduate-level calculus and linear algebra. No previous natural language processing experience required, although some background in NLP or computational linguistics will be helpful.

[Natural Language Processing with PyTorch](#) - Delip Rao 2019-01-22

Natural Language Processing (NLP) provides boundless opportunities for solving problems in artificial intelligence, making products such as Amazon Alexa and Google Translate possible. If you're a developer or data scientist new to NLP and deep learning, this practical guide shows you how to apply these methods using PyTorch, a Python-based deep learning library. Authors Delip Rao and Brian McMahon provide you with a solid grounding in NLP and deep learning algorithms and demonstrate how to use PyTorch to build applications involving rich representations of text specific to the problems you face. Each chapter includes several code examples and illustrations. Explore computational graphs and the supervised learning paradigm Master the basics of the PyTorch optimized tensor manipulation library Get an overview of traditional NLP concepts and methods Learn the basic ideas involved in building neural networks Use embeddings to represent words, sentences, documents, and other features Explore sequence prediction and generate sequence-to-sequence models Learn design patterns for building production NLP systems

Transfer Learning for Natural Language Processing - Paul Azunre 2021-08-31

Build custom NLP models in record time by adapting pre-trained machine learning models to solve specialized problems. Summary In Transfer Learning for Natural Language Processing you will learn: Fine tuning pretrained models with new domain data Picking the right model to reduce resource usage Transfer learning for neural network architectures Generating text with generative pretrained transformers Cross-lingual transfer learning with BERT Foundations for exploring NLP academic literature Training deep learning NLP models from scratch is costly, time-consuming, and requires massive amounts of data. In Transfer Learning for Natural Language Processing, DARPA researcher Paul Azunre reveals cutting-edge transfer learning techniques that apply customizable pretrained models to your own NLP architectures. You'll learn how to use transfer learning to deliver state-of-the-art results for language comprehension, even when working with limited label data. Best of all, you'll save on training time and computational costs. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Build custom NLP models in record time, even with limited datasets! Transfer learning is a machine learning technique for adapting pretrained machine learning models to solve specialized problems. This powerful approach has revolutionized natural language processing, driving improvements in machine translation, business analytics, and natural language generation. About the book Transfer Learning for Natural Language Processing teaches you to create powerful NLP solutions quickly by building on existing pretrained models. This instantly useful book provides crystal-clear explanations of the concepts you need to grok transfer learning along with hands-on examples so you can practice your new skills immediately. As you go, you'll apply state-of-the-art transfer learning methods to create a spam email

classifier, a fact checker, and more real-world applications. What's inside Fine tuning pretrained models with new domain data Picking the right model to reduce resource use Transfer learning for neural network architectures Generating text with pretrained transformers About the reader For machine learning engineers and data scientists with some experience in NLP. About the author Paul Azunre holds a PhD in Computer Science from MIT and has served as a Principal Investigator on several DARPA research programs. Table of Contents PART 1 INTRODUCTION AND OVERVIEW 1 What is transfer learning? 2 Getting started with baselines: Data preprocessing 3 Getting started with baselines: Benchmarking and optimization PART 2 SHALLOW TRANSFER LEARNING AND DEEP TRANSFER LEARNING WITH RECURRENT NEURAL NETWORKS (RNNs) 4 Shallow transfer learning for NLP 5 Preprocessing data for recurrent neural network deep transfer learning experiments 6 Deep transfer learning for NLP with recurrent neural networks PART 3 DEEP TRANSFER LEARNING WITH TRANSFORMERS AND ADAPTATION STRATEGIES 7 Deep transfer learning for NLP with the transformer and GPT 8 Deep transfer learning for NLP with BERT and multilingual BERT 9 ULMFiT and knowledge distillation adaptation strategies 10 ALBERT, adapters, and multitask adaptation strategies 11 Conclusions [Machine Learning Solutions](#) - Jalaj Thanaki 2018-04-27

Practical, hands-on solutions in Python to overcome any problem in Machine Learning Key Features Master the advanced concepts, methodologies, and use cases of machine learning Build ML applications for analytics, NLP and computer vision domains Solve the most common problems in building machine learning models Book Description Machine learning (ML) helps you find hidden insights from your data without the need for explicit programming. This book is your key to solving any kind of ML problem you might come across in your job. You'll encounter a set of simple to complex problems while building ML models, and you'll not only resolve these problems, but you'll also learn how to build projects based on each problem, with a practical approach and easy-to-follow examples. The book includes a wide range of applications: from analytics and NLP, to computer vision domains. Some of the applications you will be working on include stock price prediction, a recommendation engine, building a chat-bot, a facial expression recognition system, and many more. The problem examples we cover include identifying the right algorithm for your dataset and use cases, creating and labeling datasets, getting enough clean data to carry out processing, identifying outliers, overfitting datasets, hyperparameter tuning, and more. Here, you'll also learn to make more timely and accurate predictions. In addition, you'll deal with more advanced use cases, such as building a gaming bot, building an extractive summarization tool for medical documents, and you'll also tackle the problems faced while building an ML model. By the end of this book, you'll be able to fine-tune your models as per your needs to deliver maximum productivity. What you will learn Select the right algorithm to derive the best solution in ML domains Perform predictive analysis efficiently using ML algorithms Predict stock prices using the stock index value Perform customer analytics for an e-commerce platform Build recommendation engines for various domains Build NLP applications for the health domain Build language generation applications using different NLP techniques Build computer vision applications such as facial emotion recognition Who this book is for This book is for the intermediate users such as machine learning engineers, data engineers, data scientists, and more, who want to solve simple to complex machine learning problems in their day-to-day work and build powerful and efficient machine learning models. A basic understanding of the machine learning concepts and some experience with Python programming is all you need to get started with this book.

Natural Language Processing Recipes - Akshay Kulkarni 2019-01-29

Implement natural language processing applications with Python using a problem-solution approach. This book has numerous coding exercises that will help you to quickly deploy natural language processing techniques, such as text classification, parts of speech identification, topic modeling, text summarization, text generation, entity extraction, and sentiment analysis. Natural Language Processing Recipes starts by offering solutions for cleaning and preprocessing text data and ways to analyze it with advanced algorithms. You'll see practical applications of the semantic as well as syntactic analysis of text, as well as complex natural language processing approaches that involve text normalization, advanced preprocessing, POS tagging, and sentiment analysis. You will also learn various applications of machine learning and deep learning in natural language processing. By using the recipes in this book, you will have a toolbox of

solutions to apply to your own projects in the real world, making your development time quicker and more efficient. What You Will Learn Apply NLP techniques using Python libraries such as NLTK, TextBlob, spaCy, Stanford CoreNLP, and many more Implement the concepts of information retrieval, text summarization, sentiment analysis, and other advanced natural language processing techniques. Identify machine learning and deep learning techniques for natural language processing and natural language generation problems Who This Book Is For Data scientists who want to refresh and learn various concepts of natural language processing through coding exercises.

The Machine Learning Solutions Architect Handbook - David Ping 2022-01-21

Build highly secure and scalable machine learning platforms to support the fast-paced adoption of machine learning solutions Key Features Explore different ML tools and frameworks to solve large-scale machine learning challenges in the cloud Build an efficient data science environment for data exploration, model building, and model training Learn how to implement bias detection, privacy, and explainability in ML model development Book Description When equipped with a highly scalable machine learning (ML) platform, organizations can quickly scale the delivery of ML products for faster business value realization. There is a huge demand for skilled ML solutions architects in different industries, and this handbook will help you master the design patterns, architectural considerations, and the latest technology insights you'll need to become one. You'll start by understanding ML fundamentals and how ML can be applied to solve real-world business problems. Once you've explored a few leading problem-solving ML algorithms, this book will help you tackle data management and get the most out of ML libraries such as TensorFlow and PyTorch. Using open source technology such as Kubernetes/Kubeflow to build a data science environment and ML pipelines will be covered next, before moving on to building an enterprise ML architecture using Amazon Web Services (AWS). You'll also learn about security and governance considerations, advanced ML engineering techniques, and how to apply bias detection, explainability, and privacy in ML model development. And finally, you'll get acquainted with AWS AI services and their applications in real-world use cases. By the end of this book, you'll be able to design and build an ML platform to support common use cases and architecture patterns like a true professional. What you will learn Apply ML methodologies to solve business problems Design a practical enterprise ML platform architecture Implement MLOps for ML workflow automation Build an end-to-end data management architecture using AWS Train large-scale ML models and optimize model inference latency Create a business application using an AI service and a custom ML model Use AWS services to detect data and model bias and explain models Who this book is for This book is for data scientists, data engineers, cloud architects, and machine learning enthusiasts who want to become machine learning solutions architects. You'll need basic knowledge of the Python programming language, AWS, linear algebra, probability, and networking concepts before you get started with this handbook.

Natural Language Processing with SAS - 2020-08-31

Natural Language Processing (NLP) is a branch of artificial intelligence that helps computers understand, interpret, and emulate written or spoken human language. NLP draws from many disciplines including human-generated linguistic rules, machine learning, and deep learning to fill the gap between human communication and machine understanding. The papers included in this special collection demonstrate how NLP can be used to scale the human act of reading, organizing, and quantifying text data.

Practical Natural Language Processing with Python - Mathangi Sri 2020-12-01

Work with natural language tools and techniques to solve real-world problems. This book focuses on how natural language processing (NLP) is used in various industries. Each chapter describes the problem and solution strategy, then provides an intuitive explanation of how different algorithms work and a deeper dive on code and output in Python. Practical Natural Language Processing with Python follows a case study-based approach. Each chapter is devoted to an industry or a use case, where you address the real business problems in that industry and the various ways to solve them. You start with various types of text data before focusing on the customer service industry, the type of data available in that domain, and the common NLP problems encountered. Here you cover the bag-of-words model supervised learning technique as you try to solve the case studies. Similar depth is given to other use cases such as online reviews, bots, finance, and so on. As you cover the problems in these industries you'll also cover sentiment analysis,

named entity recognition, word2vec, word similarities, topic modeling, deep learning, and sequence to sequence modelling. By the end of the book, you will be able to handle all types of NLP problems independently. You will also be able to think in different ways to solve language problems. Code and techniques for all the problems are provided in the book. What You Will Learn Build an understanding of NLP problems in industry Gain the know-how to solve a typical NLP problem using language-based models and machine learning Discover the best methods to solve a business problem using NLP - the tried and tested ones Understand the business problems that are tough to solve Who This Book Is For Analytics and data science professionals who want to kick start NLP, and NLP professionals who want to get new ideas to solve the problems at hand.

Methods & Techniques in Deep Learning - Avik Santra 2022-12-13

Introduces multiple state-of-the-art deep learning architectures for mmwave radar in a variety of advanced applications Methods and Techniques in Deep Learning: Advancements in mmWave Radar Solutions provides a timely and authoritative overview of the use of artificial intelligence (AI)-based processing for various mmwave radar applications. Focusing on practical deep learning techniques, this comprehensive volume explains the fundamentals of deep learning, reviews cutting-edge deep metric learning techniques, describes different typologies of reinforcement learning (RL) algorithms, highlights how domain adaptation (DA) can be used for improving the performance of machine learning (ML) algorithms, and more.

Throughout the book, readers are exposed to product-ready deep learning solutions while learning skills that are relevant for building any industrial-grade, sensor-based deep learning solution. A team of authors with more than 70 filed patents and 100 published papers on AI and sensor processing illustrate how deep learning is enabling a range of advanced industrial, consumer, and automotive applications of mmwave radars. In-depth chapters cover topics including multi-modal deep learning approaches, the elemental blocks required to formulate Bayesian deep learning, how domain adaptation (DA) can be used for improving the performance of machine learning algorithms, and geometric deep learning are used for processing point clouds. In addition, the book: Discusses various advanced applications and how their respective challenges have been addressed using different deep learning architectures and algorithms Describes deep learning in the context of computer vision, natural language processing, sensor processing, and mmwave radar sensors Demonstrates how deep parametric learning reduces the number of trainable parameters and improves the data flow Presents several human-machine interface (HMI) applications such as gesture recognition, human activity classification, human localization and tracking in-cabin automotive occupancy sensing Methods and Techniques in Deep Learning: Advancements in mmWave Radar Solutions is an invaluable resource for industry professionals, researchers, and graduate students working in systems engineering, signal processing, sensors, data science and AI.

Applied Machine Learning Solutions with Python - Siddhanta Bhatta 2021-08-31

A problem-focused guide for tackling industrial machine learning issues with methods and frameworks chosen by experts. KEY FEATURES ● Popular techniques for problem formulation, data collection, and data cleaning in machine learning. ● Comprehensive and useful machine learning tools such as MLFlow, Streamlit, and many more. ● Covers numerous machine learning libraries, including Tensorflow, FastAI, Scikit-Learn, Pandas, and Numpy. DESCRIPTION This book discusses how to apply machine learning to real-world problems by utilizing real-world data. In this book, you will investigate data sources, become acquainted with data pipelines, and practice how machine learning works through numerous examples and case studies. The book begins with high-level concepts and implementation (with code!) and progresses towards the real-world of ML systems. It briefly discusses various concepts of Statistics and Linear Algebra. You will learn how to formulate a problem, collect data, build a model, and tune it. You will learn about use cases for data analytics, computer vision, and natural language processing. You will also explore nonlinear architecture, thus enabling you to build models with multiple inputs and outputs. You will get trained on creating a machine learning profile, various machine learning libraries, Statistics, and FAST API. Throughout the book, you will use Python to experiment with machine learning libraries such as Tensorflow, Scikit-learn, Spacy, and FastAI. The book will help train our models on both Kaggle and our datasets. WHAT YOU WILL LEARN ● Construct a machine learning problem, evaluate the feasibility, and gather and clean data. ● Learn to explore data first, select, and train machine learning models. ● Fine-tune

the chosen model, deploy, and monitor it in production. ● Discover popular models for data analytics, computer vision, and Natural Language Processing. ● Create a machine learning profile and contribute to the community. WHO THIS BOOK IS FOR This book caters to beginners in machine learning, software engineers, and students who want to gain a good understanding of machine learning concepts and create production-ready ML systems. This book assumes you have a beginner-level understanding of Python.

TABLE OF CONTENTS 1. Introduction to Machine Learning 2. Problem Formulation in Machine Learning 3. Data Acquisition and Cleaning 4. Exploratory Data Analysis 5. Model Building and Tuning 6. Taking Our Model into Production 7. Data Analytics Use Case 8. Building a Custom Image Classifier from Scratch 9. Building a News Summarization App Using Transformers 10. Multiple Inputs and Multiple Output Models 11. Contributing to the Community 12. Creating Your Project 13. Crash Course in Numpy, Matplotlib, and Pandas 14. Crash Course in Linear Algebra and Statistics 15. Crash Course in FastAPI

Blueprints for Text Analytics Using Python - Jens Albrecht 2021-01-12

Turning text into valuable information is essential for many businesses looking to gain a competitive advantage. There have been many improvements in natural language processing and users have a lot of options when choosing to work on a problem. However, it's not always clear which NLP tools or libraries would work for a business use--or which techniques you should use and in what order. This practical book provides theoretical background and real-world case studies with detailed code examples to help developers and data scientists obtain insight from text online. Authors Jens Albrecht, Sidharth Ramachandran, and Christian Winkler use blueprints for text-related problems that apply state-of-the-art machine learning methods in Python. If you have a fundamental understanding of statistics and machine learning along with basic programming experience in Python, you're ready to get started. You'll learn how to: Crawl and clean then explore and visualize textual data in different formats Preprocess and vectorize text for machine learning Apply methods for classification, topic analysis, summarization, and knowledge extraction Use semantic word embeddings and deep learning approaches for complex problems Work with Python NLP libraries like spaCy, NLTK, and Gensim in combination with scikit-learn, Pandas, and PyTorch

Neural Network Methods in Natural Language Processing - Yoav Goldberg 2017-04-17

Neural networks are a family of powerful machine learning models and this book focuses on their application to natural language data. The first half of the book (Parts I and II) covers the basics of supervised machine learning and feed-forward neural networks, the basics of working with machine learning over language data, and the use of vector-based rather than symbolic representations for words. It also covers the computation-graph abstraction, which allows to easily define and train arbitrary neural networks, and is the basis behind the design of contemporary neural network software libraries. The second part of the book (Parts III and IV) introduces more specialized neural network architectures, including 1D convolutional neural networks, recurrent neural networks, conditioned-generation models, and attention-based models. These architectures and techniques are the driving force behind state-of-the-art algorithms for machine translation, syntactic parsing, and many other applications. Finally, we also discuss tree-shaped networks, structured prediction, and the prospects of multi-task learning.

Deep Learning Techniques and Optimization Strategies in Big Data Analytics - Thomas, J. Joshua 2019-11-29

Many approaches have sprouted from artificial intelligence (AI) and produced major breakthroughs in the computer science and engineering industries. Deep learning is a method that is transforming the world of data and analytics. Optimization of this new approach is still unclear, however, and there's a need for research on the various applications and techniques of deep learning in the field of computing. *Deep Learning Techniques and Optimization Strategies in Big Data Analytics* is a collection of innovative research on the methods and applications of deep learning strategies in the fields of computer science and information systems. While highlighting topics including data integration, computational modeling, and scheduling systems, this book is ideally designed for engineers, IT specialists, data analysts, data scientists, engineers, researchers, academicians, and students seeking current research on deep learning methods and its application in the digital industry.