

Deep Learning Natural Language Processing In Python With Glove From Word2vec To Glove In Python And Theano Deep Learning And Natural Language Processing

Explore tools and techniques to analyze and process text with a view to building real-world NLP applications

Dependency Parsing

Natural Language Annotation for Machine Learning

Creating Neural Networks with Python

Deep Learning for NLP and Speech Recognition

Getting started with Deep Learning for Natural Language Processing

Build Intelligent Language Applications Using Deep Learning

Real-World Natural Language Processing

Deep Learning for Natural Language Processing

Speech & Language Processing

Introduction to Deep Learning

Deep Learning for Natural Language Processing

Practical Applications with Deep Learning

Learn how to build NLP applications with Deep Learning (English Edition)

Hands-On Natural Language Processing with Python

Understanding, analyzing, and generating text with Python

Natural Language Processing with PyTorch

Deep Learning for NLP and Speech Recognition

Hands-On Python Natural Language Processing

Natural Language Processing with TensorFlow

Foundations of Statistical Natural Language Processing

Python Natural Language Processing

Deep Learning for Natural Language Processing

Build Intelligent Language Applications Using Deep Learning

Practical applications with deep learning

Deep Learning for Natural Language Processing

Natural Language Processing with Python

A Machine Learning Perspective

Linguistics for the Age of AI

Solve your natural language processing problems with smart deep neural networks

A Guide to Corpus-Building for Applications

Practical Natural Language Processing

Neural Network Methods in Natural Language Processing

A Visual, Interactive Guide to Artificial Intelligence

Deep Learning in Natural Language Processing

Natural Language Processing in Action

Deep Natural Language Processing and AI Applications for Industry 5.0

Teach language to machines using Python's deep learning library

Data Science for Healthcare

Deep Learning Natural Language Processing In Python With Glove From Word2vec To Glove In Python And Theano Deep Learning And Natural Language Processing

Downloaded from archive.imba.com by guest

ALEXIA GOODMAN

Explore tools and techniques to analyze and process text with a view to building real-world NLP applications Addison-Wesley Professional

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

Dependency Parsing Packt Publishing Ltd

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.

Natural Language Annotation for Machine Learning Packt Publishing Ltd

Deep Learning for Natural Language Processing Manning Publications

Engineering Science Reference

Learn how to redesign NLP applications from scratch. KEY FEATURES • Get familiar with the basics of any Machine Learning or Deep Learning application. • Understand how does preprocessing work in NLP pipeline. • Use simple PyTorch snippets to create basic building blocks of the network commonly used in NLP. • Learn how to build a complex NLP application. • Get familiar with the advanced embedding technique, Generative network, and Audio signal processing techniques. DESCRIPTION Natural language processing (NLP) is one of the areas where many Machine Learning and Deep Learning techniques are applied. This book covers wide areas, including the fundamentals of Machine Learning, Understanding and optimizing Hyperparameters, Convolution Neural Networks (CNN), and Recurrent Neural Networks (RNN). This book not only covers the classical concept of text processing but also shares the recent advancements. This book will empower users in designing networks with the least computational and time complexity. This book not only covers basics of Natural Language Processing but also helps in deciphering the logic behind advanced concepts/architecture such as Batch Normalization, Position Embedding, DenseNet, Attention Mechanism, Highway Networks, Transformer models and Siamese Networks. This book also covers recent advancements such as ELMo-BiLM, SkipThought, and Bert. This book also covers practical implementation with step by step explanation of deep learning techniques in Topic Modelling, Text Generation, Named Entity Recognition, Text Summarization, and Language Translation. In addition to this, very advanced and open to research topics such as Generative Adversarial Network and Speech Processing are also covered. WHAT YOU WILL LEARN • Learn how to leveraging GPU for Deep Learning • Learn how to use complex embedding models such as BERT • Get familiar with the common NLP applications. • Learn how to use GANs in NLP • Learn how to process Speech data and implementing it in Speech applications WHO THIS BOOK IS FOR This book is a must-read to everyone who wishes to start the career with Machine learning and Deep Learning. This book is also for those who want to use GPU for developing Deep Learning applications. TABLE OF CONTENTS 1. Understanding the basics of learning Process 2. Text Processing Techniques 3. Representing Language

Mathematically 4. Using RNN for NLP 5. Applying CNN In NLP Tasks 6. Accelerating NLP with Advanced Embeddings 7. Applying Deep Learning to NLP tasks 8. Application of Complex Architectures in NLP 9. Understanding Generative Networks 10. Techniques of Speech Processing 11. The Road Ahead *Creating Neural Networks with Python* "O'Reilly Media, Inc." 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.

Deep Learning for NLP and Speech Recognition "O'Reilly Media, Inc."

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.

Getting started with Deep Learning for Natural Language Processing Springer

Summary: Natural Language Processing in Action is your guide to creating machines that understand human language using the power of Python with its ecosystem of packages dedicated to NLP and AI. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology: Recent advances in deep learning empower applications to understand text and speech with extreme accuracy. The result? Chatbots that can imitate real people, meaningful resume-to-job matches, superb predictive search, and automatically generated document summaries—all at a low cost. New techniques, along with accessible tools like Keras and TensorFlow, make professional-quality NLP easier than ever before. About the Book: Natural Language Processing in Action is your guide to building machines that can read and interpret human language. In it, you'll use readily available Python packages to capture the meaning in text and react accordingly. The book expands traditional NLP approaches to include neural networks, modern deep learning algorithms, and generative techniques as you tackle real-world problems like extracting dates and names, composing text, and answering free-form questions. What's inside: Some sentences in this book were written by NLP! Can you guess which ones? Working with Keras, TensorFlow, gensim, and scikit-learn. Rule-based and data-based NLP. Scalable pipelines. About the Reader: This book requires a basic understanding of deep learning and intermediate Python skills. About the Author: Hobson Lane, Cole Howard, and Hannes Max Hapke are experienced NLP engineers who use these techniques in production. Table of Contents: PART 1 - WORDY MACHINES: Packets of thought (NLP overview), Build your vocabulary (word tokenization), Math with words (TF-IDF vectors), Finding meaning in word counts (semantic analysis). PART 2 - DEEPER LEARNING (NEURAL NETWORKS): Baby steps with neural networks (perceptrons and backpropagation), Reasoning with word vectors (Word2vec), Getting words in order with convolutional neural networks (CNNs), Loopy (recurrent) neural networks (RNNs), Improving retention with long short-term memory networks, Sequence-to-sequence models and attention. PART 3 - GETTING REAL (REAL-WORLD NLP CHALLENGES): Information extraction (named entity extraction and question answering), Getting chatty (dialog engines), Scaling up (optimization, parallelization, and batch processing).

Build Intelligent Language Applications Using Deep Learning "O'Reilly Media, Inc."

Get well-versed with traditional as well as modern natural language processing concepts and techniques. Key Features: Perform various NLP tasks to build linguistic applications using Python libraries. Understand, analyze, and generate text to provide accurate results. Interpret human language using various NLP concepts, methodologies, and tools. Book Description: Natural Language Processing (NLP) is the subfield in computational linguistics that enables computers to understand, process, and analyze text. This book caters to the unmet demand for hands-on training of NLP concepts and provides exposure to real-world applications along with a solid theoretical grounding. This book starts by introducing you to the field of NLP and its applications, along with the modern Python libraries that you'll use to build your NLP-powered apps. With the help of practical examples, you'll learn how to build reasonably sophisticated NLP applications, and cover various methodologies and challenges in deploying NLP applications in the real world. You'll cover key NLP tasks such as text classification, semantic embedding, sentiment analysis, machine translation, and developing a chatbot using machine learning and deep learning techniques. The book will also help you discover how machine learning techniques play a vital role in making your linguistic apps smart. Every chapter is accompanied by examples of real-world applications to help you build impressive NLP applications of your own. By the end of this NLP book, you'll be able to work with language data, use machine learning to identify patterns in text, and get acquainted with the advancements in NLP. What you will learn: Understand how NLP powers modern applications. Explore key NLP techniques to build your natural language vocabulary. Transform text data into mathematical data structures and learn how to improve text mining models. Discover how various neural network architectures work with natural language data. Get the hang of building sophisticated text processing models using machine learning and deep learning. Check out state-of-the-art architectures that have revolutionized research in the NLP domain. Who this book is for: This NLP Python book is for anyone looking to learn NLP's theoretical and practical aspects alike. It starts with the basics and gradually covers advanced concepts to make it easy to follow for readers with varying levels of NLP proficiency. This

comprehensive guide will help you develop a thorough understanding of the NLP methodologies for building linguistic applications; however, working knowledge of Python programming language and high school level mathematics is expected.

Real-World Natural Language Processing Machine Learning Mastery

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. *Deep Learning for Natural Language Processing* Manning Publications

Gain the knowledge of various deep neural network architectures and their application areas to conquer your NLP issues. Key Features: Gain insights into the basic building blocks of natural language processing. Learn how to select the best deep neural network to solve your NLP problems. Explore convolutional and recurrent neural networks and long short-term memory networks. Book Description: Applying deep learning approaches to various NLP tasks can take your computational algorithms to a completely new level in terms of speed and accuracy. *Deep Learning for Natural Language Processing* starts off by highlighting the basic building blocks of the natural language processing domain. The book goes on to introduce the problems that you can solve using state-of-the-art neural network models. After this, delving into the various neural network architectures and their specific areas of application will help you to understand how to select the best model to suit your needs. As you advance through this deep learning book, you'll study convolutional, recurrent, and recursive neural networks, in addition to covering long short-term memory networks (LSTM). Understanding these networks will help you to implement their models using Keras. In the later chapters, you will be able to develop a trigger word detection application using NLP techniques such as attention model and beam search. By the end of this book, you will not only have sound knowledge of natural language processing but also be able to select the best text pre-processing and neural network models to solve a number of NLP issues. What you will learn: Understand various pre-processing techniques for deep learning problems. Build a vector representation of text using word2vec and GloVe. Create a named entity recognizer and parts-of-speech tagger with Apache OpenNLP. Build a machine translation model in Keras. Develop a text generation application using LSTM. Build a trigger word detection application using an attention model. Who this book is for: If you're an aspiring data scientist looking for an introduction to deep learning in the NLP domain, this is just the book for you. Strong working knowledge of Python, linear algebra, and machine learning is a must.

Speech & Language Processing IGI Global

Real-world Natural Language Processing shows you how to build the practical NLP applications that are transforming the way humans and computers work together. In Real-world Natural Language Processing you will learn how to: Design, develop, and deploy useful NLP applications. Create named entity taggers. Build machine translation systems. Construct language generation systems and chatbots. Use advanced NLP concepts such as attention and transfer learning. Real-world Natural Language Processing teaches you how to create practical NLP applications without getting bogged down in complex language theory and the mathematics of deep learning. In this engaging book, you'll explore the core tools and techniques required to build a huge range of powerful NLP apps, including chatbots, language detectors, and text classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology: Training computers to interpret and generate speech and text is a monumental challenge, and the payoff for reducing labor and improving human/computer interaction is huge! The field of Natural Language Processing (NLP) is advancing rapidly, with countless new tools and practices. This unique book offers an innovative collection of NLP techniques with applications in machine translation, voice assistants, text generation, and more. About the

book: Real-world Natural Language Processing shows you how to build the practical NLP applications that are transforming the way humans and computers work together. Guided by clear explanations of each core NLP topic, you'll create many interesting applications including a sentiment analyzer and a chatbot. Along the way, you'll use Python and open source libraries like AllenNLP and HuggingFace Transformers to speed up your development process. What's inside: Design, develop, and deploy useful NLP applications. Create named entity taggers. Build machine translation systems. Construct language generation systems and chatbots. About the reader: For Python programmers. No prior machine learning knowledge assumed. About the author: Masato Hagiwara received his computer science PhD from Nagoya University in 2009. He has interned at Google and Microsoft Research, and worked at Duolingo as a Senior Machine Learning Engineer. He now runs his own research and consulting company. Table of Contents: PART 1 BASICS: 1 Introduction to natural language processing, 2 Your first NLP application, 3 Word and document embeddings, 4 Sentence classification, 5 Sequential labeling and language modeling. PART 2 ADVANCED MODELS: 6 Sequence-to-sequence models, 7 Convolutional neural networks, 8 Attention and Transformer, 9 Transfer learning with pretrained language models. PART 3 PUTTING INTO PRODUCTION: 10 Best practices in developing NLP applications, 11 Deploying and serving NLP applications.

Introduction to Deep Learning Simon and Schuster

This book teaches you to leverage deep learning models in performing various NLP tasks along with showcasing the best practices in dealing with the NLP challenges. The book equips you with practical knowledge to implement deep learning in your linguistic applications using NLTK and Python's popular deep learning library, TensorFlow.

Deep Learning for Natural Language Processing Simon and Schuster

Become a proficient NLP data scientist by developing deep learning models for NLP and extract valuable insights from structured and unstructured data. Key Features: Get to grips with word embeddings, semantics, labeling, and high-level word representations using practical examples. Learn modern approaches to NLP and explore state-of-the-art NLP models using PyTorch. Improve your NLP applications with innovative neural networks such as RNNs, LSTMs, and CNNs. Book Description: In the internet age, where an increasing volume of text data is generated daily from social media and other platforms, being able to make sense of that data is a crucial skill. With this book, you'll learn how to extract valuable insights from text by building deep learning models for natural language processing (NLP) tasks. Starting by understanding how to install PyTorch and using CUDA to accelerate the processing speed, you'll explore how the NLP architecture works with the help of practical examples. This PyTorch NLP book will guide you through core concepts such as word embeddings, CBOW, and tokenization in PyTorch. You'll then learn techniques for processing textual data and see how deep learning can be used for NLP tasks. The book demonstrates how to implement deep learning and neural network architectures to build models that will allow you to classify and translate text and perform sentiment analysis. Finally, you'll learn how to build advanced NLP models, such as conversational chatbots. By the end of this book, you'll not only have understood the different NLP problems that can be solved using deep learning with PyTorch, but also be able to build models to solve them. What you will learn: Use NLP techniques for understanding, processing, and generating text. Understand PyTorch, its applications and how it can be used to build deep linguistic models. Explore the wide variety of deep learning architectures for NLP. Develop the skills you need to process and represent both structured and unstructured NLP data. Become well-versed with state-of-the-art technologies and exciting new developments in the NLP domain. Create chatbots using attention-based neural networks. Who this book is for: This PyTorch book is for NLP developers, machine learning and deep learning developers, and anyone interested in building intelligent language applications using both traditional NLP approaches and deep learning architectures. If you're looking to adopt modern NLP techniques and models for your development projects, this book is for you. Working knowledge of Python programming, along with basic working knowledge of NLP tasks, is required.

Practical Applications with Deep Learning Packt Publishing Ltd

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.

Learn how to build NLP applications with Deep Learning (English Edition) Packt Publishing Ltd

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

Hands-On Natural Language Processing with Python

Pearson Education India

Neural networks are a family of powerful machine learning models. This book focuses on the application of neural network models 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. [Understanding, analyzing, and generating text with Python](#) Simon

and Schuster

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 PyTorch Simon and Schuster 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.

Deep Learning for NLP and Speech Recognition Cambridge University Press

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.

Hands-On Python Natural Language Processing Apress

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by sitting down and writing programs," the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to Tensorflow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.

Related with Deep Learning Natural Language Processing In Python With Glove From Word2vec To Glove In Python And Theano Deep Learning And Natural Language Processing:

- Black Jeopardy Questions And Answers : [click here](#)