
Mastering Natural Language Processing With Python

AI Unleashed: A Holistic Guide to Mastering Artificial Intelligence
Natural Language Processing Fundamentals for Developers
Mastering Machine Learning Algorithms
Mastering Machine Learning: A Comprehensive Guide to Success
Natural Language Processing
Artificial Intelligence, Real Profits
Essential Math for AI
Mastering Your Prompt Engineering Super Power
Mastering Image Creation with Midjourney AI Tools
Natural Language Processing with Python Quick Start Guide
Mastering Artificial Intelligence and Machine Learning
Mastering Predictive Analytics with scikit-learn and TensorFlow
Hands-On Natural Language Processing with Python
Mastering Machine Learning with Python in Six Steps
Mastering Text Mining with R
Mastering Natural Language Processing with Python
Python Machine Learning Blueprints
Mastering Predictive Analytics with R
Learning Path
Natural Language Processing: Python and NLTK
Introduction to Natural Language Processing
Slavonic Natural Language Processing in the 21st Century
Mastering Machine Learning Algorithms
Mastering Machine Learning with Core ML and Python
Mastering Transformers
Highlights in Practical Applications of Agents, Multi-Agent Systems, and Trust-worthiness. The PAAMS Collection
Context in Computing
Mastering ChatGPT
Mastering the Art of Artificial Intelligence
Mastering Machine Learning with R
Mastering PyTorch
Foundations of Statistical Natural Language Processing
Mastering Machine Learning with R - Third Edition
Transformers for Natural Language Processing
MASTERING PYTORCH -
Artificial Intelligence, Real Profits
Mastering Data Science
Mastering spaCy

MACK MARISSA

AI Unleashed: A Holistic Guide to Mastering Artificial Intelligence

Tribun EU

This book constitutes the refereed proceedings of the workshops co-located with the 18th International Conference on Practical Applications of Agents and Multi-Agent Systems, PAAMS 2020, held in L'Aquila, Italy, in October 2020. The total of 21 full and 13 short papers presented in this volume were carefully reviewed and selected from 57 submissions. The papers in this volume stem from the following meetings: Workshop on Agent-Based Artificial Markets Computational Economics (ABAM); Workshop on Agents and Edge-AI (AgEdAI); Workshop on Character Computing (C2); Workshop on MAS for Complex Networks and Social Computation (CNSC); Workshop on Decision Support, Recommendation, and Persuasion in Artificial Intelligence (DeRePAI); Workshop on Multi-Agent Systems and Simulation (MAS&S); Workshop on Multi-agent based Applications for Energy Markets, Smart Grids and Sustainable Energy Systems (MASGES); Workshop on Smart Cities and Intelligent Agents (SCIA).

[Natural Language Processing Fundamentals for Developers](#)
AppCoda

Welcome to "Mastering Machine Learning: A Comprehensive Guide to Success." In this book, we embark on an exciting journey into the world of machine learning (ML), exploring its concepts, techniques, and practical applications. Whether you are a beginner taking your first steps into the field or an experienced practitioner seeking to deepen your knowledge, this comprehensive guide will equip you with the tools, strategies, and insights needed to succeed in the ever-evolving landscape of ML. Machine learning is a rapidly advancing field that has revolutionized industries and transformed the way we tackle complex problems. From personalized recommendations and speech recognition systems to autonomous vehicles and medical diagnostics, machine learning has become an integral part of our daily lives. Its ability to analyze vast amounts of data, identify

patterns, and make predictions has paved the way for groundbreaking advancements across various domains. However, mastering machine learning requires more than just understanding the algorithms and techniques. It requires a holistic approach that encompasses data collection and preparation, exploratory data analysis, model building, evaluation, deployment, and continuous learning. It also demands a deep understanding of the ethical and social implications of machine learning, ensuring responsible and fair use of this powerful technology. In this book, we have carefully crafted 20 comprehensive chapters that cover a wide range of topics, from the fundamentals of machine learning to advanced techniques and future trends. Each chapter provides a deep dive into a specific aspect of machine learning, offering tips, recommendations, and strategies for success. You will learn about various algorithms, data preprocessing techniques, model evaluation methods, interpretability approaches, and much more. Throughout the book, we emphasize a practical approach to machine learning. Real-world examples, case studies, and hands-on exercises are incorporated to help you gain a deeper understanding of the concepts and apply them to your own projects. We believe that active learning and practical experience are crucial for mastering machine learning, and we encourage you to explore, experiment, and build your own models. While this book serves as a comprehensive guide, it is important to note that machine learning is a rapidly evolving field. New algorithms, techniques, and technologies are constantly emerging, and staying up-to-date with the latest advancements is essential. However, the principles and foundations discussed in this book will provide you with a solid framework to adapt and navigate the ever-changing landscape of machine learning. Whether you are an aspiring data scientist, a software engineer, a researcher, or a business professional, this book is designed to be your trusted companion in your journey to mastering machine learning. By the time you reach the end, you will have gained a deep understanding of the fundamental concepts, acquired practical skills for applying machine learning in real-world scenarios, and developed the mindset needed to tackle complex challenges and

drive innovation. Get ready to embark on an exciting adventure into the world of machine learning. Let's begin our journey towards mastering machine learning and unlocking its full potential. Happy learning!

[Mastering Machine Learning Algorithms](#) Packt Publishing Ltd
SpaCy, a fast, user-friendly library for teaching computers to understand text, simplifies NLP techniques, such as speech tagging and syntactic dependencies, so you can easily extract information, attributes, and objects from massive amounts of text to then document, measure, and analyze. This Learning Path is a hands-on introduction to using SpaCy to discover insights through natural language processing. While end-to-end natural language processing solutions can be complex, you'll learn the linguistics, algorithms, and machine learning skills to get the job done.
[Mastering Machine Learning: A Comprehensive Guide to Success](#)
Packt Publishing Ltd

Learn advanced techniques to improve the performance and quality of your predictive models
Key Features
Use ensemble methods to improve the performance of predictive analytics models
Implement feature selection, dimensionality reduction, and cross-validation techniques
Develop neural network models and master the basics of deep learning
Book Description
Python is a programming language that provides a wide range of features that can be used in the field of data science. Mastering Predictive Analytics with scikit-learn and TensorFlow covers various implementations of ensemble methods, how they are used with real-world datasets, and how they improve prediction accuracy in classification and regression problems. This book starts with ensemble methods and their features. You will see that scikit-learn provides tools for choosing hyperparameters for models. As you make your way through the book, you will cover the nitty-gritty of predictive analytics and explore its features and characteristics. You will also be introduced to artificial neural networks and TensorFlow, and how it is used to create neural networks. In the final chapter, you will explore factors such as computational power, along with improvement methods and software enhancements for efficient predictive analytics. By the end of this book, you will be well-versed in using deep neural

networks to solve common problems in big data analysis. What you will learn Use ensemble algorithms to obtain accurate predictions Apply dimensionality reduction techniques to combine features and build better models Choose the optimal hyperparameters using cross-validation Implement different techniques to solve current challenges in the predictive analytics domain Understand various elements of deep neural network (DNN) models Implement neural networks to solve both classification and regression problems Who this book is for Mastering Predictive Analytics with scikit-learn and TensorFlow is for data analysts, software engineers, and machine learning developers who are interested in implementing advanced predictive analytics using Python. Business intelligence experts will also find this book indispensable as it will teach them how to progress from basic predictive models to building advanced models and producing more accurate predictions. Prior knowledge of Python and familiarity with predictive analytics concepts are assumed.

Natural Language Processing Apress

Maximize your NLP capabilities while creating amazing NLP projects in Python About This Book* Learn to implement various NLP tasks in Python* Gain insights into the current and budding research topics of NLP* This is a comprehensive step-by-step guide to help students and researchers create their own projects based on real-life applications Who This Book Is For This book is for intermediate level developers in NLP with a reasonable knowledge level and understanding of Python. What You Will Learn* Implement string matching algorithms and normalization techniques* Implement statistical language modeling techniques* Get an insight into developing a stemmer, lemmatizer, morphological analyzer, and morphological generator* Develop a search engine and implement POS tagging concepts and statistical modeling concepts involving the n gram approach* Familiarize yourself with concepts such as the Treebank construct, CFG construction, the CYK Chart Parsing algorithm, and the Earley Chart Parsing algorithm* Develop an NER-based system and understand and apply the concepts of sentiment analysis* Understand and implement the concepts of Information Retrieval and text summarization* Develop a Discourse Analysis System and Anaphora Resolution based system In Detail Natural Language Processing is one of the fields of computational

linguistics and artificial intelligence that is concerned with human-computer interaction. It provides a seamless interaction between computers and human beings and gives computers the ability to understand human speech with the help of machine learning. This book will give you expertise on how to employ various NLP tasks in Python, giving you an insight into the best practices when designing and building NLP-based applications using Python. It will help you become an expert in no time and assist you in creating your own NLP projects using NLTK. You will sequentially be guided through applying machine learning tools to develop various models. We'll give you clarity on how to create training data and how to implement major NLP applications such as Named Entity Recognition, Question Answering System, Discourse Analysis, Transliteration, Word Sense disambiguation, Information Retrieval, Sentiment Analysis, Text Summarization, and Anaphora Resolution.

Artificial Intelligence, Real Profits Cybellium Ltd

Mastering Natural Language Processing with Python

Essential Math for AI Mercury Learning and Information

This book is for developers who are looking for an overview of basic concepts in Natural Language Processing. It casts a wide net of techniques to help developers who have a range of technical backgrounds. Numerous code samples and listings are included to support myriad topics. The first chapter shows you various details of managing data that are relevant for NLP. The next pair of chapters contain NLP concepts, followed by another pair of chapters with Python code samples to illustrate those NLP concepts. Chapter 6 explores applications, e.g., sentiment analysis, recommender systems, COVID-19 analysis, spam detection, and a short discussion regarding chatbots. The final chapter presents the Transformer architecture, BERT-based models, and the GPT family of models, all of which were developed during the past three years and considered SOTA ("state of the art"). The appendices contain introductory material (including Python code samples) on regular expressions and probability/statistical concepts. Companion files with source code and figures are included. FEATURES: Covers extensive topics related to natural language processing Includes separate appendices on regular expressions and probability/statistics Features companion files with source code and figures from the book. The companion files are available online by emailing the

publisher with proof of purchase at info@merclearning.com. *Mastering Your Prompt Engineering Super Power* Packt Publishing Ltd

Updated and revised second edition of the bestselling guide to exploring and mastering the most important algorithms for solving complex machine learning problems Key Features Updated to include new algorithms and techniques Code updated to Python 3.8 & TensorFlow 2.x New coverage of regression analysis, time series analysis, deep learning models, and cutting-edge applications Book Description Mastering Machine Learning Algorithms, Second Edition helps you harness the real power of machine learning algorithms in order to implement smarter ways of meeting today's overwhelming data needs. This newly updated and revised guide will help you master algorithms used widely in semi-supervised learning, reinforcement learning, supervised learning, and unsupervised learning domains. You will use all the modern libraries from the Python ecosystem – including NumPy and Keras – to extract features from varied complexities of data. Ranging from Bayesian models to the Markov chain Monte Carlo algorithm to Hidden Markov models, this machine learning book teaches you how to extract features from your dataset, perform complex dimensionality reduction, and train supervised and semi-supervised models by making use of Python-based libraries such as scikit-learn. You will also discover practical applications for complex techniques such as maximum likelihood estimation, Hebbian learning, and ensemble learning, and how to use TensorFlow 2.x to train effective deep neural networks. By the end of this book, you will be ready to implement and solve end-to-end machine learning problems and use case scenarios. What you will learn Understand the characteristics of a machine learning algorithm Implement algorithms from supervised, semi-supervised, unsupervised, and RL domains Learn how regression works in time-series analysis and risk prediction Create, model, and train complex probabilistic models Cluster high-dimensional data and evaluate model accuracy Discover how artificial neural networks work – train, optimize, and validate them Work with autoencoders, Hebbian networks, and GANs Who this book is for This book is for data science professionals who want to delve into complex ML algorithms to understand how various machine learning models can be built. Knowledge of Python programming is required.

Mastering Image Creation with Midjourney AI Tools Rick Spair

This volume explores how context has been and can be used in computing to model human behaviors, actions and communications as well as to manage data and knowledge. It addresses context management and exploitation of context for sharing experience across domains. The book serves as a user-centric guide for readers wishing to develop context-based applications, as well as an intellectual reference on the concept of context. It provides a broad yet deep treatment of context in computing and related areas that depend heavily on computing. The coverage is broad because of its cross-disciplinary nature but treats topics at a sufficient depth to permit a reader to implement context in his/her computational endeavors. The volume addresses how context can be integrated in software and systems and how it can be used in a computing environment. Furthermore, the use of context to represent the human dimension, individually as well as collectively is explained. Contributions also include descriptions of how context has been represented in formal as well as non-formal, structured approaches. The last section describes several human behavior representation paradigms based on the concept of context as its central representational element. The depth and breadth of this content is certain to provide useful as well as intellectually enriching information to readers of diverse backgrounds who have an interest in or are intrigued by using context to assist in their representation of the real world.

Natural Language Processing with Python Quick Start Guide Packt Publishing Ltd

"Natural Language Processing is one of the fields of computational linguistics and artificial intelligence that is concerned with human-computer interaction. It provides a seamless interaction between computers and human beings and gives computers the ability to understand human speech with the help of machine learning. This course will give you expertise on how to employ various NLP tasks in Python, giving you an insight into the best practices when designing and building NLP-based applications using Python. It will help you become an expert in no time and assist you in creating your own NLP projects using NLTK. You will sequentially be guided through applying machine learning tools to develop various models. We'll give you clarity on how to create training data and how to implement major NLP

applications such as Named Entity Recognition, Question Answering System, Discourse Analysis, Transliteration, Word Sense disambiguation, Information Retrieval, Text Summarization, and Anaphora Resolution."--Resource description page.

Mastering Artificial Intelligence and Machine 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.

Mastering Predictive Analytics with scikit-learn and TensorFlow MIT Press

R offers a free and open source environment that is perfect for both learning and deploying predictive modeling solutions in the

real world. With its constantly growing community and plethora of packages, R offers the functionality to deal with a truly vast array of problems. This book is designed to be both a guide and a reference for moving beyond the basics of predictive modeling. The book begins with a dedicated chapter on the language of models and the predictive modeling process. Each subsequent chapter tackles a particular type of model, such as neural networks, and focuses on the three important questions of how the model works, how to use R to train it, and how to measure and assess its performance using real world data sets. By the end of this book, you will have explored and tested the most popular modeling techniques in use on real world data sets and mastered a diverse range of techniques in predictive analytics.

Hands-On Natural Language Processing with Python Packt Publishing Ltd

This textbook presents an up-to-date and comprehensive overview of Natural Language Processing (NLP), from basic concepts to core algorithms and key applications. Further, it contains seven step-by-step NLP workshops (total length: 14 hours) offering hands-on practice with essential Python tools like NLTK, spaCy, TensorFlow Keras, Transformer and BERT. The objective of this book is to provide readers with a fundamental grasp of NLP and its core technologies, and to enable them to build their own NLP applications (e.g. Chatbot systems) using Python-based NLP tools. It is both a textbook and NLP tool-book intended for the following readers: undergraduate students from various disciplines who want to learn NLP; lecturers and tutors who want to teach courses or tutorials for undergraduate/graduate students on NLP and related AI topics; and readers with various backgrounds who want to learn NLP, and more importantly, to build workable NLP applications after completing its 14 hours of Python-based workshops.

Mastering Machine Learning with Python in Six Steps Packt Publishing Ltd

Statistical approaches to processing natural language text have become dominant in recent years. This foundational text is the first comprehensive introduction to statistical natural language processing (NLP) to appear. The book contains all the theory and algorithms needed for building NLP tools. It provides broad but rigorous coverage of mathematical and linguistic foundations, as well as detailed discussion of statistical methods, allowing

students and researchers to construct their own implementations. The book covers collocation finding, word sense disambiguation, probabilistic parsing, information retrieval, and other applications.

Mastering Text Mining with R Packt Publishing Ltd

The book reflects new advances in Slavonic natural language processing at the beginning of the 21st century. The whole book was dedicated to Karel Pala.

Mastering Natural Language Processing with Python Athena Publishing

Master advanced techniques and algorithms for deep learning with PyTorch using real-world examples Key Features Understand how to use PyTorch 1.x to build advanced neural network models Learn to perform a wide range of tasks by implementing deep learning algorithms and techniques Gain expertise in domains such as computer vision, NLP, Deep RL, Explainable AI, and much more Book Description Deep learning is driving the AI revolution, and PyTorch is making it easier than ever before for anyone to build deep learning applications. This PyTorch book will help you uncover expert techniques to get the most out of your data and build complex neural network models. The book starts with a quick overview of PyTorch and explores using convolutional neural network (CNN) architectures for image classification. You'll then work with recurrent neural network (RNN) architectures and transformers for sentiment analysis. As you advance, you'll apply deep learning across different domains, such as music, text, and image generation using generative models and explore the world of generative adversarial networks (GANs). You'll not only build and train your own deep reinforcement learning models in PyTorch but also deploy PyTorch models to production using expert tips and techniques. Finally, you'll get to grips with training large models efficiently in a distributed manner, searching neural architectures effectively with AutoML, and rapidly prototyping models using PyTorch and fast.ai. By the end of this PyTorch book, you'll be able to perform complex deep learning tasks using PyTorch to build smart artificial intelligence models. What you will learn Implement text and music generating models using PyTorch Build a deep Q-network (DQN) model in PyTorch Export universal PyTorch models using Open Neural Network Exchange (ONNX) Become well-versed with rapid prototyping using PyTorch with fast.ai Perform neural architecture search effectively using AutoML Easily interpret machine learning (ML) models written in

PyTorch using Captum Design ResNets, LSTMs, Transformers, and more using PyTorch Find out how to use PyTorch for distributed training using the torch.distributed API Who this book is for This book is for data scientists, machine learning researchers, and deep learning practitioners looking to implement advanced deep learning paradigms using PyTorch 1.x. Working knowledge of deep learning with Python programming is required.

Python Machine Learning Blueprints Packt Publishing

In a world driven by data and powered by artificial intelligence, there's a superpower that's changing the game: Prompt Engineering. Join Diana Ashcroft, a seasoned data scientist and educator, on a journey through the dynamic landscape of prompt engineering in her latest book, *Mastering Your Prompt Engineering Super Power*. Prompt engineering is the key to unlocking the full potential of AI. In *Mastering Your Prompt Engineering Super Power*, Diana Ashcroft delves into the heart of this transformative field and reveals its immense significance. You'll discover how prompt engineering is reshaping industries, powering innovation, and shaping the future of society. Whether you're a seasoned AI professional or just starting your journey, *Mastering Your Prompt Engineering Super Power* is your guide to mastering prompt engineering. Diana takes complex concepts and distills them into practical, down-to-earth knowledge that anyone can grasp. You'll explore the realms of Natural Language Processing (NLP), Computer Vision, and more, gaining the skills needed to harness prompt engineering's incredible potential. Prompt engineering isn't just a buzzword; it's a force that's driving change in every sector. Diana provides real-world examples of how prompt engineering is making waves in industries like healthcare, finance, e-commerce, and beyond. You'll see how AI-powered prompts are enhancing productivity, improving customer experiences, and even revolutionizing education. *Mastering Your Prompt Engineering Super Power* isn't just a book; it's your passport to becoming a prompt engineering master. Diana guides you through hands-on techniques, tools, and frameworks used by professionals in the field. You'll learn to wield the power of AI-driven prompts to tackle complex tasks, from data preprocessing to model optimization. As we stand on the precipice of a new era, Diana Ashcroft illuminates the path forward. Discover how prompt engineering is shaping the future, from enabling smarter virtual assistants to aiding legal professionals in document analysis. The

possibilities are endless, and *Mastering Your Prompt Engineering Super Power* equips you to seize them.

Mastering Predictive Analytics with R Cybellium Ltd

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.

[Learning Path](#) Springer

Discover a project-based approach to mastering machine learning concepts by applying them to everyday problems using libraries such as scikit-learn, TensorFlow, and Keras. Key Features: Get to grips with Python's machine learning libraries including scikit-learn, TensorFlow, and Keras. Implement advanced concepts and popular machine learning algorithms in real-world projects. Build analytics, computer vision, and neural network projects. Book Description: Machine learning is transforming the way we understand and interact with the world around us. This book is the perfect guide for you to put your knowledge and skills into practice and use the Python ecosystem to cover key domains in machine learning. This second edition covers a range of libraries from the Python ecosystem, including TensorFlow and Keras, to help you implement real-world machine learning projects. The book begins by giving you an overview of machine learning with Python. With the help of complex datasets and optimized techniques, you'll go on to understand how to apply advanced concepts and popular machine learning algorithms to real-world projects. Next, you'll cover projects from domains such as predictive analytics to analyze the stock market and recommendation systems for GitHub repositories. In addition to this, you'll also work on projects from the NLP domain to create a custom news feed using frameworks such as scikit-learn, TensorFlow, and Keras. Following this, you'll learn how to build an advanced chatbot, and scale things up using PySpark. In the concluding chapters, you can look forward to exciting insights into

deep learning and you'll even create an application using computer vision and neural networks. By the end of this book, you'll be able to analyze data seamlessly and make a powerful impact through your projects. What you will learn: Understand the Python data science stack and commonly used algorithms. Build a model to forecast the performance of an Initial Public Offering (IPO) over an initial discrete trading window. Understand NLP concepts by creating a custom news feed. Create applications that will recommend GitHub repositories based on ones you've starred, watched, or forked. Gain the skills to build a chatbot from scratch using PySpark. Develop a market-prediction app using stock data. Delve into advanced concepts such as computer vision, neural networks, and deep learning. Who this book is for: This book is for machine learning practitioners, data scientists, and deep learning enthusiasts who want to take their machine learning skills to the next level by building real-world projects. The intermediate-level guide will help you to implement libraries from the Python ecosystem to build a variety of projects addressing various machine learning domains. Knowledge of Python programming and machine learning concepts will be helpful. *Natural Language Processing: Python and NLTK* Packt Publishing Ltd

Build end-to-end industrial-strength NLP models using advanced morphological and syntactic features in spaCy to create real-world applications with ease. Key Features: Gain an overview of what spaCy offers for natural language processing. Learn details of spaCy's features and how to use them effectively. Work through practical recipes using spaCy. Book Description: spaCy is an industrial-grade, efficient NLP Python library. It offers various pre-trained models and ready-to-use features. Mastering spaCy provides you with end-to-end coverage of spaCy's features and real-world applications. You'll begin by installing spaCy and downloading models, before progressing to spaCy's features and

prototyping real-world NLP apps. Next, you'll get familiar with visualizing with spaCy's popular visualizer displaCy. The book also equips you with practical illustrations for pattern matching and helps you advance into the world of semantics with word vectors. Statistical information extraction methods are also explained in detail. Later, you'll cover an interactive business case study that shows you how to combine all spaCy features for creating a real-world NLP pipeline. You'll implement ML models such as sentiment analysis, intent recognition, and context resolution. The book further focuses on classification with popular frameworks such as TensorFlow's Keras API together with spaCy. You'll cover popular topics, including intent classification and sentiment analysis, and use them on popular datasets and interpret the classification results. By the end of this book, you'll be able to confidently use spaCy, including its linguistic features, word vectors, and classifiers, to create your own NLP apps. What you will learn: Install spaCy, get started easily, and write your first Python script. Understand core linguistic operations of spaCy. Discover how to combine rule-based components with spaCy statistical models. Become well-versed with named entity and keyword extraction. Build your own ML pipelines using spaCy. Apply all the knowledge you've gained to design a chatbot using spaCy. Who this book is for: This book is for data scientists and machine learners who want to excel in NLP as well as NLP developers who want to master spaCy and build applications with it. Language and speech professionals who want to get hands-on with Python and spaCy and software developers who want to quickly prototype applications with spaCy will also find this book helpful. Beginner-level knowledge of the Python programming language is required to get the most out of this book. A beginner-level understanding of linguistics such as parsing, POS tags, and semantic similarity will also be useful.

Related with Mastering Natural Language Processing With Python:

© [Mastering Natural Language Processing With Python Whats Wrong In Spanish Language](#)

© [Mastering Natural Language Processing With Python What Was The Role Of Lineage Groups In African Society](#)

© [Mastering Natural Language Processing With Python Whats Analysis Paralysis](#)