
Mastering Natural Language Processing With Python

Natural Language Processing In 5 Minutes | What Is NLP And How Does It Work? | Simplilearn
Neuro Linguistic Programming Techniques You Can Use Instantly
What Is The Best NLP Book? Mastering Tokenization in NLP | Natural Language Processing | NLP | Python | Tutorial 02 Natural Language Processing with spaCy
Python - Course for Beginners
NLP Secrets: What are the big secrets to mastering NLP? How I'd Learn NLP in 2024 (If I Had to Start Over)
Mastering Natural Language Processing (NLP): A 5-Hour Deep Dive from Basics to Advanced Projects
How I'd learn ML in 2024 (if I could start over)
NATURAL LANGUAGE PROCESSING With Python | Theory
Hands-On Exercise Natural Language Processing In 10 Minutes | NLP Tutorial For Beginners | NLP Training | Simplilearn

Real-World Natural Language Processing

Build effective real-world NLP applications using NER, RNNs, seq2seq models, Transformers, and more

Foundations of Statistical Natural Language Processing

Mastering Java Machine Learning

Going from a Python developer to an effective Natural Language Processing Engineer

Mastering Natural Language Processing with Python

Build innovative deep neural network architectures for NLP with Python, PyTorch, TensorFlow, BERT, RoBERTa, and more

Mastering Machine Learning with Python in Six Steps

A practical guide to text analysis with Python, Gensim, spaCy, and Keras

Mastering Azure Machine Learning

A comprehensive guide to understanding machine learning and developing AI-based apps for iOS.

Mastering Machine Learning with scikit-learn

Multilingual Natural Language Processing Applications

Natural Language Processing with Python and spaCy

Develop Deep Learning Models for your Natural Language Problems

A Practical Implementation Guide to Predictive Data Analytics Using Python

Transfer Learning for Natural Language Processing
Mastering Machine Learning with Core ML and Python
Build powerful neural network architectures using advanced PyTorch 1.x features

*Mastering Natural
Language Processing
With Python*

*OMB No.
0638184296459 edited
by*

ARIANA WENDY

REAL-WORLD NATURAL LANGUAGE PROCESSING

Machine Learning Mastery
Acquire and analyze data from all corners
of the social web with Python About This
Book Make sense of highly unstructured
social media data with the help of the
insightful use cases provided in this guide
Use this easy-to-follow, step-by-step guide
to apply analytics to complicated and
messy social data This is your one-stop
solution to fetching, storing, analyzing,
and visualizing social media data Who This
Book Is For This book is for intermediate
Python developers who want to engage
with the use of public APIs to collect data
from social media platforms and perform
statistical analysis in order to produce
useful insights from data. The book

assumes a basic understanding of the
Python Standard Library and provides
practical examples to guide you toward
the creation of your data analysis project
based on social data. What You Will Learn
Interact with a social media platform via
their public API with Python Store social
data in a convenient format for data
analysis Slice and dice social data using
Python tools for data science Apply text
analytics techniques to understand what
people are talking about on social media
Apply advanced statistical and analytical
techniques to produce useful insights from
data Build beautiful visualizations with
web technologies to explore data and
present data products In Detail Your social
media is filled with a wealth of hidden data
- unlock it with the power of Python.
Transform your understanding of your
clients and customers when you use
Python to solve the problems of
understanding consumer behavior and
turning raw data into actionable customer
insights. This book will help you acquire

and analyze data from leading social
media sites. It will show you how to
employ scientific Python tools to mine
popular social websites such as Facebook,
Twitter, Quora, and more. Explore the
Python libraries used for social media
mining, and get the tips, tricks, and insider
insight you need to make the most of
them. Discover how to develop data
mining tools that use a social media API,
and how to create your own data analysis
projects using Python for clear insight from
your social data. Style and approach This
practical, hands-on guide will help you
learn everything you need to perform data
mining for social media. Throughout the
book, we take an example-oriented
approach to use Python for data analysis
and provide useful tips and tricks that you
can use in day-to-day tasks.

**BUILD EFFECTIVE REAL-WORLD NLP
APPLICATIONS USING NER, RNNs,**

SEQ2SEQ MODELS, TRANSFORMERS, AND MORE

Packt Publishing Ltd

Explore fundamental to advanced Python 3 topics in six steps, all designed to make you a worthy practitioner. This updated version's approach is based on the "six degrees of separation" theory, which states that everyone and everything is a maximum of six steps away and presents each topic in two parts: theoretical concepts and practical implementation using suitable Python 3 packages. You'll start with the fundamentals of Python 3 programming language, machine learning history, evolution, and the system development frameworks. Key data mining/analysis concepts, such as exploratory analysis, feature dimension reduction, regressions, time series forecasting and their efficient implementation in Scikit-learn are covered as well. You'll also learn commonly used model diagnostic and tuning techniques. These include optimal probability cutoff point for class creation, variance, bias, bagging, boosting, ensemble voting, grid search, random search, Bayesian

optimization, and the noise reduction technique for IoT data. Finally, you'll review advanced text mining techniques, recommender systems, neural networks, deep learning, reinforcement learning techniques and their implementation. All the code presented in the book will be available in the form of iPython notebooks to enable you to try out these examples and extend them to your advantage. What You'll Learn Understand machine learning development and frameworks Assess model diagnosis and tuning in machine learning Examine text mining, natural language processing (NLP), and recommender systems Review reinforcement learning and CNN Who This Book Is For Python developers, data engineers, and machine learning engineers looking to expand their knowledge or career into machine learning area.

FOUNDATIONS OF STATISTICAL NATURAL LANGUAGE PROCESSING

Packt Publishing Ltd

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.

MASTERING JAVA MACHINE LEARNING

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.

GOING FROM A PYTHON DEVELOPER TO AN EFFECTIVE NATURAL LANGUAGE PROCESSING ENGINEER

AppCoda

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.

MASTERING NATURAL LANGUAGE PROCESSING WITH PYTHON

Mastering Natural Language Processing with Python

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.

Build innovative deep neural network architectures for NLP with Python, PyTorch, TensorFlow, BERT, RoBERTa, and more 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.

Mastering Machine Learning with Python in Six Steps Packt Publishing Ltd

Take a problem-solving approach to learning all about transformers and get up and running in no time by implementing methodologies that will build the future of NLP

Key Features

- Explore quick prototyping with up-to-date Python libraries to create effective solutions to industrial problems
- Solve advanced NLP problems such as named-entity recognition, information extraction, language generation, and conversational AI
- Monitor your model's performance with the help of BertViz, exBERT, and TensorBoard

Book Description

Transformer-based language models have dominated natural language processing (NLP) studies and have now become a new paradigm. With this book, you'll learn how to build various transformer-based NLP applications using the Python

Transformers library. The book gives you an introduction to Transformers by showing you how to write your first hello-world program. You'll then learn how a tokenizer works and how to train your own tokenizer. As you advance, you'll explore the architecture of autoencoding models, such as BERT, and autoregressive models, such as GPT. You'll see how to train and fine-tune models for a variety of natural language understanding (NLU) and natural language generation (NLG) problems, including text classification, token classification, and text representation. This book also helps you to learn efficient models for challenging problems, such as long-context NLP tasks with limited computational capacity. You'll also work with multilingual and cross-lingual problems, optimize models by monitoring their performance, and discover how to deconstruct these models for interpretability and explainability. Finally, you'll be able to deploy your transformer models in a production environment. By the end of this NLP book, you'll have learned how to use Transformers to solve advanced NLP problems using advanced models. What you will learn

Explore state-

of-the-art NLP solutions with the Transformers library Train a language model in any language with any transformer architecture Fine-tune a pre-trained language model to perform several downstream tasks Select the right framework for the training, evaluation, and production of an end-to-end solution Get hands-on experience in using TensorBoard and Weights & Biases Visualize the internal representation of transformer models for interpretability Who this book is for This book is for deep learning researchers, hands-on NLP practitioners, as well as ML/NLP educators and students who want to start their journey with Transformers. Beginner-level machine learning knowledge and a good command of Python will help you get the best out of this book.

[A practical guide to text analysis with Python, Gensim, spaCy, and Keras](#)

"O'Reilly Media, Inc."

Since their introduction in 2017, transformers have quickly become the dominant architecture for achieving state-of-the-art results on a variety of natural language processing tasks. If you're a data scientist or coder, this practical book

shows you how to train and scale these large models using Hugging Face Transformers, a Python-based deep learning library. Transformers have been used to write realistic news stories, improve Google Search queries, and even create chatbots that tell corny jokes. In this guide, authors Lewis Tunstall, Leandro von Werra, and Thomas Wolf, among the creators of Hugging Face Transformers, use a hands-on approach to teach you how transformers work and how to integrate them in your applications. You'll quickly learn a variety of tasks they can help you solve. Build, debug, and optimize transformer models for core NLP tasks, such as text classification, named entity recognition, and question answering Learn how transformers can be used for cross-lingual transfer learning Apply transformers in real-world scenarios where labeled data is scarce Make transformer models efficient for deployment using techniques such as distillation, pruning, and quantization Train transformers from scratch and learn how to scale to multiple GPUs and distributed environments

MASTERING AZURE MACHINE LEARNING

CRC Press

Build and deploy intelligent applications for natural language processing with Python by using industry standard tools and recently popular methods in deep learning Key Features A no-math, code-driven programmer's guide to text processing and NLP Get state of the art results with modern tooling across linguistics, text vectors and machine learning Fundamentals of NLP methods from spaCy, gensim, scikit-learn and PyTorch Book Description NLP in Python is among the most sought after skills among data scientists. With code and relevant case studies, this book will show how you can use industry-grade tools to implement NLP programs capable of learning from relevant data. We will explore many modern methods ranging from spaCy to word vectors that have reinvented NLP. The book takes you from the basics of NLP to building text processing applications. We start with an introduction to the basic vocabulary along with a workflow for building NLP applications. We use

industry-grade NLP tools for cleaning and pre-processing text, automatic question and answer generation using linguistics, text embedding, text classifier, and building a chatbot. With each project, you will learn a new concept of NLP. You will learn about entity recognition, part of speech tagging and dependency parsing for Q and A. We use text embedding for both clustering documents and making chatbots, and then build classifiers using scikit-learn. We conclude by deploying these models as REST APIs with Flask. By the end, you will be confident building NLP applications, and know exactly what to look for when approaching new challenges. What you will learn

Understand classical linguistics in using English grammar for automatically generating questions and answers from a free text corpus

Work with text embedding models for dense number representations of words, subwords and characters in the English language for exploring document clustering

Deep Learning in NLP using PyTorch with a code-driven introduction to PyTorch

Using an NLP project management Framework for estimating timelines and organizing your project into

stages

Hack and build a simple chatbot application in 30 minutes

Deploy an NLP or machine learning application using Flask as RESTFUL APIs

Who this book is for

Programmers who wish to build systems that can interpret language. Exposure to Python programming is required. Familiarity with NLP or machine learning vocabulary will be helpful, but not mandatory.

A comprehensive guide to understanding machine learning and developing AI-based apps for iOS.

Apress

Deep learning methods are achieving state-of-the-art results on challenging machine learning problems such as describing photos and translating text from one language to another. In this new laser-focused Ebook, finally cut through the math, research papers and patchwork descriptions about natural language processing. Using clear explanations, standard Python libraries and step-by-step tutorial lessons you will discover what natural language processing is, the promise of deep learning in the field, how to clean and prepare text data for modeling, and how to develop deep

learning models for your own natural language processing projects.

Mastering Machine Learning with scikit-learn Packt Publishing

This book provides hands-on training in NLP tools and techniques with intrinsic details. Apart from gaining expertise, you will be able to carry out novel state-of-the-art research using the skills gained.

Multilingual Natural Language Processing Applications Packt Publishing Ltd

Master text-taming techniques and build effective text-processing applications with R

About This Book

Develop all the relevant skills for building text-mining apps with R with this easy-to-follow guide

Gain in-depth understanding of the text mining process with lucid implementation in the R language

Example-rich guide that lets you gain high-quality information from text data

Who This Book Is For

If you are an R programmer, analyst, or data scientist who wants to gain experience in performing text data mining and analytics with R, then this book is for you. Exposure to working with statistical methods and language processing would be helpful.

What You Will Learn

Get acquainted with some of the highly efficient R packages

such as OpenNLP and RWeka to perform various steps in the text mining process. Access and manipulate data from different sources such as JSON and HTTP. Process text using regular expressions. Get to know the different approaches of tagging texts, such as POS tagging, to get started with text analysis. Explore different dimensionality reduction techniques, such as Principal Component Analysis (PCA), and understand its implementation in R. Discover the underlying themes or topics that are present in an unstructured collection of documents, using common topic models such as Latent Dirichlet Allocation (LDA). Build a baseline sentence-completing application. Perform entity extraction and named entity recognition using R. In Detail Text Mining (or text data mining or text analytics) is the process of extracting useful and high-quality information from text by devising patterns and trends. R provides an extensive ecosystem to mine text through its many frameworks and packages. Starting with basic information about the statistics concepts used in text mining, this book will teach you how to access, cleanse, and process text using the R language and will

equip you with the tools and the associated knowledge about different tagging, chunking, and entailment approaches and their usage in natural language processing. Moving on, this book will teach you different dimensionality reduction techniques and their implementation in R. Next, we will cover pattern recognition in text data utilizing classification mechanisms, perform entity recognition, and develop an ontology learning framework. By the end of the book, you will develop a practical application from the concepts learned, and will understand how text mining can be leveraged to analyze the massively available data on social media. Style and approach This book takes a hands-on, example-driven approach to the text mining process with lucid implementation in R.

NATURAL LANGUAGE PROCESSING WITH PYTHON AND SPACY

Simon and Schuster

An introduction to natural language processing with Python using spaCy, a leading Python natural language processing library. Natural Language

Processing with Python and spaCy will show you how to create NLP applications like chatbots, text-condensing scripts, and order-processing tools quickly and easily. You'll learn how to leverage the spaCy library to extract meaning from text intelligently; how to determine the relationships between words in a sentence (syntactic dependency parsing); identify nouns, verbs, and other parts of speech (part-of-speech tagging); and sort proper nouns into categories like people, organizations, and locations (named entity recognizing). You'll even learn how to transform statements into questions to keep a conversation going. You'll also learn how to:

- Work with word vectors to mathematically find words with similar meanings (Chapter 5)
- Identify patterns within data using spaCy's built-in displaCy visualizer (Chapter 7)
- Automatically extract keywords from user input and store them in a relational database (Chapter 9)
- Deploy a chatbot app to interact with users over the internet (Chapter 11)

"Try This" sections in each chapter encourage you to practice what you've learned by expanding the book's example scripts to handle a wider range of

inputs, add error handling, and build professional-quality applications. By the end of the book, you'll be creating your own NLP applications with Python and spaCy.

Develop Deep Learning Models for your Natural Language Problems Packt Publishing Ltd

This practical book provides 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, this second edition includes code updated for Python 3, shows you how to scale up for larger data sets, and covers the semantic web. 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 [A Practical Implementation Guide to Predictive Data Analytics Using Python](#) Packt Publishing Ltd

Machine learning, now more than ever, plays a pivotal role in almost everything we do in our digital lives. Whether it's interacting with a virtual assistant like Siri or typing out a message to a friend, machine learning is the technology facilitating those actions. It's clear that machine learning is here to stay, and as such, it's a vital skill to have in the upcoming decades. This book covers Core ML in-depth. You will learn how to create and deploy your own machine learning model. On top of that, you will learn about Turi Create, Create ML, Keras, Firebase, and Jupyter Notebooks, just to name a few. These are a few examples of professional tools which are staples for many machine learning experts. By going through this book, you'll also become proficient with Python, the language that's

most frequently used for machine learning. Plus, you would have created a handful of ready-to-use apps such as barcode scanners, image classifiers, and language translators. Most importantly, you will master the ins-and-outs of Core ML.

Transfer Learning for Natural Language Processing Packt Publishing Ltd

The application of deep learning methods to problems in natural language processing has generated significant progress across a wide range of natural language processing tasks. For some of these applications, deep learning models now approach or surpass human performance. While the success of this approach has transformed the engineering methods of machine learning in artificial intelligence, the significance of these achievements for the modelling of human learning and representation remains unclear. *Deep Learning and Linguistic Representation* looks at the application of a variety of deep learning systems to several cognitively interesting NLP tasks. It also considers the extent to which this work illuminates our understanding of the way in which humans acquire and

represent linguistic knowledge. Key Features: combines an introduction to deep learning in AI and NLP with current research on Deep Neural Networks in computational linguistics. is self-contained and suitable for teaching in computer science, AI, and cognitive science courses; it does not assume extensive technical training in these areas. provides a compact guide to work on state of the art systems that are producing a revolution across a range of difficult natural language tasks.

Mastering Machine Learning with Core ML and Python Apress

Become an expert at building and deploying enterprise-grade data applications in Java About This Book* This comprehensive book shows you exactly how you can take your Java data science applications to production seamlessly* Dive deep into analytics, supervised and unsupervised learning, and much more with ease* Explore Java's various libraries to efficiently build and deploy data applications for the enterprise Who This Book Is For This book is for those Java developers who are comfortable with developing applications in Java and are

familiar with the basic concepts of data science. This is the go-to book for anyone looking to master the subject using Java. If you are willing to build efficient data applications in your enterprise environment without changing your existing stack, this book is for you! What you will learn* Get a solid understanding of the data processing toolbox available in Java* Explore the data science ecosystem available in Java and other JVM languages* Understand when to use Java and what is best to do outside of Java* Deal with the machine learning task at hand and bring the results directly to production* Get state-of-the-art performance with xgboost and deeplearning4j* Build applications that scale and process large amounts of data in real time In Detail Java is the language of choice if you want to bring data science to production, thanks to its stability and rich set of libraries. Major big data solutions including Hadoop are written in Java. This book will teach you how to perform data analysis on big data in a much more sophisticated manner. If you are willing to take your data products to enterprise without changing your stack, this book will tell you how to do it with

ease. This book will quickly brush up on what you already know about using Java in data science applications and will then dive quickly into the advanced concepts to implement data science in production. The book covers topics such as advanced data science algorithms, preparing tricky data, advanced clustering, regression, classification, prediction, machine learning, and more. We'll teach you how data science can be used effectively to analyze unstructured data and big data. This book will enable you to tackle the problems of advanced visualization, advanced statistics, scaling data science applications, deploying these applications in production, and many more. You will also learn about natural language processing, real-time analytics, deep learning, and neural networks.

Build powerful neural network architectures using advanced PyTorch 1.x features "O'Reilly Media, Inc."

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 SMART, AI-DRIVEN LINGUISTIC APPLICATIONS USING DEEP LEARNING AND NLP TECHNIQUES

Simon and Schuster

Become an advanced practitioner with this progressive set of master classes on application-oriented machine learning About This Book Comprehensive coverage of key topics in machine learning with an emphasis on both the theoretical and practical aspects More than 15 open source Java tools in a wide range of techniques, with code and practical usage. More than 10 real-world case studies in machine learning highlighting techniques ranging from data ingestion up to analyzing the results of experiments, all preparing the user for the practical, real-world use of tools and data analysis. Who This Book Is For This book will appeal to anyone with a serious interest in topics in Data Science or those already working in related areas: ideally, intermediate-level data analysts and data scientists with experience in Java. Preferably, you will have experience with the fundamentals of machine learning and now have a desire to explore the area further, are up to

grappling with the mathematical complexities of its algorithms, and you wish to learn the complete ins and outs of practical machine learning. What You Will Learn Master key Java machine learning libraries, and what kind of problem each can solve, with theory and practical guidance. Explore powerful techniques in each major category of machine learning such as classification, clustering, anomaly detection, graph modeling, and text mining. Apply machine learning to real-world data with methodologies, processes, applications, and analysis. Techniques and experiments developed around the latest specializations in machine learning, such as deep learning, stream data mining, and active and semi-supervised learning. Build high-performing, real-time, adaptive predictive models for batch- and stream-

based big data learning using the latest tools and methodologies. Get a deeper understanding of technologies leading towards a more powerful AI applicable in various domains such as Security, Financial Crime, Internet of Things, social networking, and so on. In Detail Java is one of the main languages used by practicing data scientists; much of the Hadoop ecosystem is Java-based, and it is certainly the language that most production systems in Data Science are written in. If you know Java, Mastering Machine Learning with Java is your next step on the path to becoming an advanced practitioner in Data Science. This book aims to introduce you to an array of advanced techniques in machine learning, including classification, clustering,

anomaly detection, stream learning, active learning, semi-supervised learning, probabilistic graph modeling, text mining, deep learning, and big data batch and stream machine learning. Accompanying each chapter are illustrative examples and real-world case studies that show how to apply the newly learned techniques using sound methodologies and the best Java-based tools available today. On completing this book, you will have an understanding of the tools and techniques for building powerful machine learning models to solve data science problems in just about any domain. Style and approach A practical guide to help you explore machine learning—and an array of Java-based tools and frameworks—with the help of practical examples and real-world use cases.

Related with Mastering Natural Language Processing With Python:

[© Mastering Natural Language Processing With Python How To Change Commentary Language In Fifa 22](#)

[© Mastering Natural Language Processing With Python How To Become A Fashion Historian](#)

[© Mastering Natural Language Processing With Python How To Become A Financial Historian](#)