

Text Analytics with Python

A Practitioner's Guide to Natural Language Processing

Leverage Natural Language Processing (NLP) in Python and learn how to set up your own robust environment for performing text analytics. This second edition has gone through a major revamp and introduces several significant changes and new topics based on the recent trends in NLP.

You'll see how to use the latest state-of-the-art frameworks in NLP, coupled with machine learning and deep learning models for supervised sentiment analysis powered by Python to solve actual case studies. Start by reviewing Python for NLP fundamentals on strings and text data and move on to engineering representation methods for text data, including both traditional statistical models and newer deep learning-based embedding models. Improved techniques and new methods around parsing and processing text are discussed as well.

Text summarization and topic models have been overhauled so the book showcases how to build, tune, and interpret topic models in the context of an interest dataset on NIPS conference papers. Additionally, the book covers text similarity techniques with a real-world example of movie recommenders, along with sentiment analysis using supervised and unsupervised techniques.

There is also a chapter dedicated to semantic analysis where you'll see how to build your own named entity recognition (NER) system from scratch. While the overall structure of the book remains the same, the entire code base, modules, and chapters has been updated to the latest Python 3.x release.

What You'll Learn


- Understand NLP and text syntax, semantics and structure
- Discover text cleaning and feature engineering
- Review text classification and text clustering
- Assess text summarization and topic models
- Study deep learning for NLP

Who This Book Is For

IT professionals, data analysts, developers, linguistic experts, data scientists and engineers and basically anyone with a keen interest in linguistics, analytics and generating insights from textual data.

Leverage Natural Language Processing (NLP) in Python and learn how to set up your own robust environment for performing text analytics. This second edition has gone through a major revamp and introduces several significant changes and new topics based on the recent trends in NLP. You'll see how to use the latest state-of-the-art frameworks in NLP, coupled with machine learning and deep learning models for supervised sentiment analysis powered by Python to solve actual case studies. Start by reviewing Python for NLP fundamentals on strings and text data and move on to engineering representation methods for text data, including both traditional statistical models and newer deep learning-based embedding models. Improved techniques and new methods around parsing and processing text are discussed as well. Text summarization and topic models have been overhauled so the book showcases how to build, tune, and interpret topic models in the context of an interest dataset on NIPS conference papers. Additionally, the book covers text similarity techniques with a real-world example of movie recommenders, along with sentiment analysis using supervised and unsupervised techniques. There is also a chapter dedicated to semantic analysis where you'll see how to build your own named entity recognition (NER) system from scratch. While the overall structure of the book remains the same, the entire code base, modules, and chapters has been updated to the latest Python 3.x release.

What You'll Learn



48,14 €

44,99 € (zzgl. MwSt.)

Nicht mehr lieferbar

Artikelnummer: 9781484243534

Medium: Buch

ISBN: 978-1-4842-4353-4

Verlag: Apress

Erscheinungstermin: 22.05.2019

Sprache(n): Englisch

Auflage: 2. Auflage 2019

Produktform: Kartoniert

Gewicht: 1293 g

Seiten: 674

Format (B x H): 178 x 254 mm

