



# Python Natural Language Processing Cookbook : Over 60 Recipes for Building Powerful NLP Solutions Using Python and LLM Libraries /

Antic, Zhenya,  
author

Monografía

Harness the power of Natural Language Processing to overcome real-world text analysis challenges with this recipe-based roadmap written by two seasoned NLP experts with vast experience transforming various industries with their NLP prowess. You'll be able to make the most of the latest NLP advancements, including large language models (LLMs), and leverage their capabilities through Hugging Face transformers. Through a series of hands-on recipes, you'll master essential techniques such as extracting entities and visualizing text data. The authors will expertly guide you through building pipelines for sentiment analysis, topic modeling, and question-answering using popular libraries like spaCy, Gensim, and NLTK. You'll also learn to implement RAG pipelines to draw out precise answers from a text corpus using LLMs. This second edition expands your skillset with new chapters on cutting-edge LLMs like GPT-4, Natural Language Understanding (NLU), and Explainable AI (XAI)-fostering trust and transparency in your NLP models. By the end of this book, you'll be equipped with the skills to apply advanced text processing techniques, use pre-trained transformer models, build custom NLP pipelines to extract valuable insights from text data to drive informed decision-making

<https://rebiunoda.pro.baratznet.cloud:28443/OpacDiscovery/public/catalog/detail/b2FpOmNlbgVcmF0aW9uOmVzLmJhcmF0ei5yZW4vMzYzNjkyNDU>

**Título:** Python Natural Language Processing Cookbook Over 60 Recipes for Building Powerful NLP Solutions Using Python and LLM Libraries Zhenya Antic and Saurabh Chakravarty

**Edición:** Second edition

**Editorial:** Birmingham, England Packt Publishing Ltd. [2024] 2024

**Descripción física:** 1 online resource (312 pages)

**Nota general:** Includes index

**Contenido:** Cover -- Title page -- Copyright and credits -- Dedication -- Contributors -- Table of Contents -- Preface -- Chapter 1: Learning NLP Basics -- Technical requirements -- Dividing text into sentences -- Getting ready -- How to do it... -- There's more... -- See also -- Dividing sentences into words - tokenization -- Getting ready -- How to do it -- There's more... -- See also -- Part of speech tagging -- Getting ready -- How to do it... --

There's more -- There's more -- See also -- Combining similar words - lemmatization -- Getting ready -- How to do it... -- There's more... -- Removing stopwords -- Getting ready -- How to do it... -- There's more... -- Chapter 2: Playing with Grammar -- Technical requirements -- Counting nouns - plural and singular nouns -- Getting ready -- How to do it... -- There's more... -- Getting the dependency parse -- Getting ready -- How to do it... -- See also -- Extracting noun chunks -- Getting ready -- How to do it... -- There's more... -- See also -- Extracting subjects and objects of the sentence -- Getting ready -- How to do it... -- There's more... -- Finding patterns in text using grammatical information -- Getting ready -- How to do it... -- See also -- Chapter 3: Representing Text - Capturing Semantics -- Technical requirements -- Creating a simple classifier -- Getting ready -- How to do it... -- There's more... -- Putting documents into a bag of words -- Getting ready -- How to do it... -- Constructing an N-gram model -- Getting ready -- How to do it... -- There's more... -- Representing texts with TF-IDF -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Using word embeddings -- Getting ready -- How to do it... -- There's more... -- See also -- Training your own embeddings model -- Getting ready -- How to do it... -- There's more... -- See also -- Using BERT and OpenAI embeddings instead of word embeddings -- Getting ready -- How to do it... -- There's more... -- See also -- Retrieval augmented generation (RAG) -- Getting ready -- How to do it... -- Chapter 4: Classifying Texts -- Technical requirements -- Getting the dataset and evaluation ready -- Getting ready -- How to do it... -- Performing rule-based text classification using keywords -- Getting ready -- How to do it... -- Clustering sentences using K-Means - unsupervised text classification -- Getting ready -- How to do it... -- Using SVMs for supervised text classification -- Getting ready -- How to do it... -- There's more... -- Training a spaCy model for supervised text classification -- Getting ready -- How to do it... -- Classifying texts using OpenAI models -- Getting ready -- How to do it... -- Chapter 5: Getting Started with Information Extraction -- Technical requirements -- Using regular expressions -- Getting ready -- How to do it... -- There's more... -- Finding similar strings - Levenshtein distance -- Getting ready -- How to do it... -- There's more... -- Extracting keywords -- Getting ready -- How to do it... -- There's more... -- Performing named entity recognition using spaCy -- Getting ready -- How to do it... -- There's more... -- Training your own NER model with spaCy -- Getting ready -- How to do it... -- See also -- Fine-tuning BERT for NER -- Getting ready -- How to do it... -- Chapter 6: Topic Modeling -- Technical requirements -- LDA topic modeling with gensim -- Getting ready -- How to do it... -- There's more... -- Community detection clustering with SBERT -- Getting ready -- How to do it... -- K-Means topic modeling with BERT -- Getting ready -- How to do it... -- Topic modeling using BERTopic -- Getting ready -- How to do it... -- There's more... -- Using contextualized topic models -- Getting ready -- How to do it... -- See also -- Chapter 7: Visualizing Text Data -- Technical requirements -- Visualizing the dependency parse -- Getting ready -- How to do it... -- Visualizing parts of speech -- Getting ready -- How to do it... -- Visualizing NER -- Getting ready -- How to do it... -- Creating a confusion matrix plot -- Getting ready -- How to do it... -- Constructing word clouds -- Getting ready -- How to do it... -- There's more... -- See also -- Visualizing topics from Gensim -- Getting ready -- How to do it... -- See also -- Visualizing topics from BERTopic -- Getting ready -- How to do it... -- See also -- Chapter 8: Transformers and Their Applications -- Technical requirements -- Loading a dataset -- Getting ready -- How to do it... -- Tokenizing the text in your dataset -- Getting ready -- How to do it... -- Classifying text -- Getting ready -- How to do it... -- Using a zero-shot classifier -- Getting ready -- How to do it... -- Generating text -- Getting ready -- How to do it... -- There's more... -- Language translation -- Getting ready -- How to do it... -- Chapter 9: Natural Language Understanding -- Technical requirements -- Answering questions from a short text passage -- Getting ready -- How to do it... -- Answering questions from a long text passage -- Getting ready -- How to do it... -- See also -- Answering questions from a document corpus in an extractive manner -- Getting ready -- How to do it... -- See also -- Answering questions from a document corpus in an abstractive manner -- Getting ready -- How to do it... -- See also -- Summarizing text using pre-trained models based on Transformers -- Getting ready -- How to do it... -- There's more... -- See also -- Detecting sentence entailment -- Getting ready -- How to do it... -- There's more... -- Enhancing explainability via a classifier-invariant approach -- Getting ready -- How to do it... -- There's more... -- Enhancing explainability via text generation -- Getting ready -- How to do it... -- Chapter 10: Generative AI and Large Language Models -- Technical requirements -- Model access -- Running an LLM locally -- Getting ready -- How to do it... -- Running an LLM to follow instructions -- Getting ready -- How to do it... -- There's more... -- Augmenting an LLM with external data -- Executing a simple prompt-to-LLM chain -- Augmenting the LLM with external content -- Creating a chatbot using an LLM -- Getting ready -- How to do it... -- Generating code using an LLM -- Getting ready -- How to do it... -- There's more... -- Generating a SQL query using human-defined requirements -- Getting ready -- How to do it... -- Agents - making an LLM to reason and act -- Getting ready -- How to do it... -- Using OpenAI models instead of local ones -- Index -- Other Books You May Enjoy

**ISBN:** 1-80324-144-6

**Materia:** Python (Computer program language) Natural language processing (Computer science)

**Autores:** Chakravarty, Saurabh, author

**Enlace a formato físico adicional:** 1-80324-574-3

---

## **Baratz Innovación Documental**

- Gran Vía, 59 28013 Madrid
- (+34) 91 456 03 60
- [informa@baratz.es](mailto:informa@baratz.es)