

Python for Natural Language Processing

Programming with NumPy, scikit-learn, Keras, and PyTorch


Since the last edition of this book (2014), progress has been astonishing in all areas of Natural Language Processing, with recent achievements in Text Generation that spurred a media interest going beyond the traditional academic circles. Text Processing has meanwhile become a mainstream industrial tool that is used, to various extents, by countless companies. As such, a revision of this book was deemed necessary to catch up with the recent breakthroughs, and the author discusses models and architectures that have been instrumental in the recent progress of Natural Language Processing.

As in the first two editions, the intention is to expose the reader to the theories used in Natural Language Processing, and to programming examples that are essential for a deep understanding of the concepts. Although present in the previous two editions, Machine Learning is now even more pregnant, having replaced many of the earlier techniques to process text. Many new techniques build on the availability of text.

Using Python notebooks, the reader will be able to load small corpora, format text, apply the models through executing pieces of code, gradually discover the theoretical parts by possibly modifying the code or the parameters, and traverse theories and concrete problems through a constant interaction between the user and the machine. The data sizes and hardware requirements are kept to a reasonable minimum so that a user can see instantly, or at least quickly, the results of most experiments on most machines.

The book does not assume a deep knowledge of Python, and an introduction to this language aimed at Text Processing is given in Ch. 2, which will enable the reader to touch all the programming concepts, including NumPy arrays and PyTorch tensors as fundamental structures to represent and process numerical data in Python, or Keras for training Neural Networks to classify texts. Covering topics like Word Segmentation and Part-of-Speech and Sequence Annotation, the textbook also gives an in-depth overview of Transformers (for instance, BERT), Self-Attention and Sequence-to-Sequence Architectures.

Since the last edition of this book (2014), progress has been astonishing in all areas of Natural Language Processing, with recent achievements in Text Generation that spurred a media interest going beyond the traditional academic circles. Text Processing has meanwhile become a mainstream industrial tool that is used, to various extents, by countless companies. As such, a revision of this book was deemed necessary to catch up with the recent breakthroughs, and the author discusses models and architectures that have been instrumental in the recent progress of Natural Language Processing. As in the first two editions, the intention is to expose the reader to the theories used in Natural Language Processing, and to programming examples that are essential for a deep understanding of the concepts. Although present in the previous two editions, Machine Learning is now even more pregnant, having replaced many of the earlier techniques to process text. Many new techniques build on the availability of text. Using Python notebooks, the reader will be able to load small corpora, format text, apply the models through executing pieces of code, gradually discover the theoretical parts by possibly modifying the code or the parameters, and traverse theories and concrete problems through a constant interaction between the user and the machine. The data sizes and hardware requirements are kept to a reasonable minimum so that a user can see instantly, or at least quickly, the results of most experiments on most machines. The book does not assume a deep knowledge of Python, and an introduction to this language aimed at Text Processing is given in Ch. 2, which will enable the reader to touch all the programming concepts, including NumPy arrays and PyTorch tensors as fundamental structures to represent and process numerical data in Python, or Keras for training Neural Networks to classify texts. Covering topics like Word Segmentation and Part-of-



64,19 €

59,99 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9783031575488

Medium: Buch

ISBN: 978-3-031-57548-8

Verlag: Springer

Erscheinungstermin: 10.07.2024

Sprache(n): Englisch

Auflage: Third Auflage 2024

Serie: Cognitive Technologies

Produktform: Gebunden

Gewicht: 978 g

Seiten: 520

Format (B x H): 160 x 241 mm

