


Deep Learning for Natural Language Processing

Deep Learning is becoming increasingly important in a technology-dominated world. However, the building of computational models that accurately represent linguistic structures is complex, as it involves an in-depth knowledge of neural networks, and the understanding of advanced mathematical concepts such as calculus and statistics. This book makes these complexities accessible to those from a humanities and social sciences background, by providing a clear introduction to deep learning for natural language processing. It covers both theoretical and practical aspects, and assumes minimal knowledge of machine learning, explaining the theory behind natural language in an easy-to-read way. It includes pseudo code for the simpler algorithms discussed, and actual Python code for the more complicated architectures, using modern deep learning libraries such as PyTorch and Hugging Face. Providing the necessary theoretical foundation and practical tools, this book will enable readers to immediately begin building real-world, practical natural language processing systems.



125,70 €

117,48 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9781316515662

Medium: Buch

ISBN: 978-1-316-51566-2

Verlag: Cambridge University Press

Erscheinungstermin: 12.01.2024

Sprache(n): Englisch

Auflage: Erscheinungsjahr 2024

Produktform: Gebunden

Gewicht: 652 g

Seiten: 342

Format (B x H): 157 x 235 mm

