

Introduction to Deep Learning


From Logical Calculus to Artificial Intelligence

This textbook presents a concise, accessible and engaging first introduction to deep learning, offering a wide range of connectionist models which represent the current state-of-the-art. The text explores the most popular algorithms and architectures in a simple and intuitive style, explaining the mathematical derivations in a step-by-step manner. The content coverage includes convolutional networks, LSTMs, Word2vec, RBMs, DBNs, neural Turing machines, memory networks and autoencoders. Numerous examples in working Python code are provided throughout the book, and the code is also supplied separately at an accompanying website.

Topics and features: introduces the fundamentals of machine learning, and the mathematical and computational prerequisites for deep learning; discusses feed-forward neural networks, and explores the modifications to these which can be applied to any neural network; examines convolutional neural networks, and the recurrent connections to a feed-forward neural network; describes the notion of distributed representations, the concept of the autoencoder, and the ideas behind language processing with deep learning; presents a brief history of artificial intelligence and neural networks, and reviews interesting open research problems in deep learning and connectionism.

This clearly written and lively primer on deep learning is essential reading for graduate and advanced undergraduate students of computer science, cognitive science and mathematics, as well as fields such as linguistics, logic, philosophy, and psychology.

This textbook presents a concise, accessible and engaging first introduction to deep learning, offering a wide range of connectionist models which represent the current state-of-the-art. The text explores the most popular algorithms and architectures in a simple and intuitive style, explaining the mathematical derivations in a step-by-step manner. The content coverage includes convolutional networks, LSTMs, Word2vec, RBMs, DBNs, neural Turing machines, memory networks and autoencoders. Numerous examples in working Python code are provided throughout the book, and the code is also supplied separately at an accompanying website. Topics and features: introduces the fundamentals of machine learning, and the mathematical and computational prerequisites for deep learning; discusses feed-forward neural networks, and explores the modifications to these which can be applied to any neural network; examines convolutional neural networks, and the recurrent connections to a feed-forward neural network; describes the notion of distributed representations, the concept of the autoencoder, and the ideas behind language processing with deep learning; presents a brief history of artificial intelligence and neural networks, and reviews interesting open research problems in deep learning and connectionism. This clearly written and lively primer on deep learning is essential reading for graduate and advanced undergraduate students of computer science, cognitive science and mathematics, as well as fields such as linguistics, logic, philosophy, and psychology.



58,84 €

54,99 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9783319730035

Medium: Buch

ISBN: 978-3-319-73003-5

Verlag: Springer

Erscheinungstermin: 15.02.2018

Sprache(n): Englisch

Auflage: 1. Auflage 2018

Serie: Undergraduate Topics in Computer Science

Produktform: Kartoniert

Gewicht: 324 g

Seiten: 191

Format (B x H): 155 x 235 mm

