

Introduction to Deep Learning

Neural Networks, Large Language Models and Agentic AI

This textbook introduces deep learning in a style that is accessible, rigorous, and grounded in working code. It walks through the most widely used algorithms and architectures step by step, with mathematical derivations kept intuitive and Python examples woven through every chapter.

The **second edition** keeps everything from the first, including convolutional networks, LSTMs, Word2vec, RBMs, DBNs, neural Turing machines, memory networks, and autoencoders. It then covers the systems that have reshaped the field since: generative adversarial networks, the transformer architecture and its attention mechanism, the full training pipeline behind modern large language models (LLMs), prompt engineering with real-life guardrail scenarios, parameter-efficient fine-tuning with LoRA, retrieval-augmented generation with vector databases, knowledge graphs, and agentic AI systems illustrated through an industrial case study.


Topics and features:

- Introduces fundamentals of machine learning and mathematical and computational prerequisites for deep learning
- Discusses feed-forward neural networks, convolutional networks, and recurrent architectures, and explores the modifications applicable to any neural network
- Covers the transformer architecture from first principles, including self-attention, multi-head attention, positional encoding, and a minimal annotated implementation
- Reviews open research problems, from hallucinations and quadratic scaling to alignment faking and the interpretability of model internals

This proven, textbook is written for graduate and advanced undergraduate students of computer science, cognitive science, and mathematics. It should prove equally valuable for readers in linguistics, logic, philosophy, and psychology.

Sandro Skansi is an Associate Professor at the University of Zagreb, Croatia, where he teaches logic, political philosophy, artificial intelligence, and cognitive science. **Kristina Šekrst** is a research associate at the University of Zagreb and a principal engineer at Preamble AI.

This textbook introduces deep learning in a style that is accessible, rigorous, and grounded in working code. It walks through the most widely used algorithms and architectures step by step, with mathematical derivations kept intuitive and Python examples woven through every chapter. The second edition keeps everything from the first, including convolutional networks, LSTMs, Word2vec, RBMs, DBNs, neural Turing machines, memory networks, and autoencoders. It then covers the systems that have reshaped the field since: generative adversarial networks, the transformer architecture and its attention mechanism, the full training pipeline behind modern large language models (LLMs), prompt engineering with real-life guardrail scenarios, parameter-efficient fine-tuning with LoRA, retrieval-augmented generation with vector databases, knowledge graphs, and agentic AI systems illustrated through an industrial case study. Topics and features:- Introduces fundamentals of machine learning and mathematical and computational prerequisites for deep learning- Discusses feed-forward neural networks, convolutional networks, and recurrent architectures, and explores the modifications applicable to any neural network- Covers the transformer architecture from first principles, including self-attention, multi-head attention, positional encoding, and a minimal annotated implementation- Reviews open research problems, from hallucinations and quadratic scaling to alignment faking and the interpretability of model internals This proven, fully revised textbook is written for graduate and advanced undergraduate students of computer science, cognitive science, and mathematics. It should prove



53,49 €
49,99 € (zzgl. MwSt.)

vorbestellbar, Erscheinungstermin ca. Oktober 2026

Artikelnummer: 9783032254597
Medium: Buch
ISBN: 978-3-032-25459-7
Verlag: Springer Nature Switzerland AG
Erscheinungstermin: 28.10.2026
Sprache(n): Englisch
Auflage: 2. Auflage 2027
Serie: Undergraduate Topics in Computer Science
Produktform: Kartoniert
Seiten: 104
Format (B x H): 155 x 235 mm

