


Advanced Data Science and Analytics with Python

The second edition of Advanced Data Science and Analytics with Python reflects the rapid transformation of artificial intelligence in recent years. While preserving its practical, modular structure, this edition significantly expands coverage of the techniques shaping modern AI practice. The deep learning chapter has been substantially broadened to include reinforcement learning and generative adversarial networks, alongside a fully developed exploration of transformer architectures. Generative AI now takes centre stage, with dedicated coverage of self-attention, BERT, GPT, large language model evaluation and API-based interaction. Emerging agentic systems are introduced as part of the evolving AI landscape. Natural language processing has been enhanced with word embeddings, contextual representations and vector search, while network analysis now includes graph representation learning and embedding techniques. The chapter on data product deployment has been strengthened with modern Core ML workflows and new coverage of on-device Foundation Models, bridging experimentation and production. Fully updated for the contemporary Python ecosystem, this edition equips practitioners with the tools and architectural understanding required to design, build and deploy intelligent systems in today's AI-driven world. Jesús Rogel-Salazar is a lead data scientist, founder of RogueLoop, working for companies such as Ortus Technologies, Tynpa Health Technologies, Barclays, AKQA, IBM Data Science Studio and Dow Jones. He is a visiting researcher at the Imperial College Business School, UK and a member of the School of Physics, Astronomy and Mathematics at the University of Hertfordshire, UK.

The second edition of Advanced Data Science and Analytics with Python reflects the rapid transformation of artificial intelligence in recent years. While preserving its practical, modular structure, this edition significantly expands coverage of the techniques shaping modern AI practice. The deep learning chapter has been substantially broadened to include reinforcement learning and generative adversarial networks, alongside a fully developed exploration of transformer architectures. Generative AI now takes centre stage, with dedicated coverage of self-attention, BERT, GPT, large language model evaluation and API-based interaction. Emerging agentic systems are introduced as part of the evolving AI landscape. Natural language processing has been enhanced with word embeddings, contextual representations and vector search, while network analysis now includes graph representation learning and embedding techniques. The chapter on data product deployment has been strengthened with modern Core ML workflows and new coverage of on-device Foundation Models, bridging experimentation and production. Fully updated for the contemporary Python ecosystem, this edition equips practitioners with the tools and architectural understanding required to design, build and deploy intelligent systems in today's AI-driven world.



185,50 €

185,50 € (zzgl. MwSt.)

vorbestellbar, Erscheinungstermin ca. August 2026

Artikelnummer: 9781041103608

Medium: Buch

ISBN: 978-1-041-10360-8

Verlag: Taylor & Francis Ltd

Erscheinungstermin: 25.08.2026

Sprache(n): Englisch

Auflage: 2. Auflage 2026

Serie: Chapman & Hall/CRC Data Mining and Knowledge Discovery Series

Produktform: Gebunden

Gewicht: 1330 g

Seiten: 632

Format (B x H): 191 x 235 mm

