


Artificial Intelligence and Safety

A Practical Guide for Programmers and Decision Makers

In a world where machines learn, evolve, and sometimes outthink us, Artificial Intelligence and Safety offers a rare, integrated lens into both the power and peril of intelligent systems. Drawing on the authors' unique blend of programming mastery and enterprise risk leadership, this guide is your compass in the age of intelligent machines. This book goes beyond traditional risk management frameworks by integrating the latest advancements in artificial intelligence (AI) and generative AI. Combining real-world case studies, Socratic questioning, and the wisdom of AI governance pioneers, it journeys through the evolution of AI, the rise of deep learning, and the practical implications of large language models (LLMs) and retrieval-augmented generation (RAG). The authors translate complex technical knowledge into actionable insights and teach how to use Python not just for prediction, but for protection—turning code into a risk management ally. Readers will understand how to embed ethics by design, assess emerging risks, and confidently apply AI governance frameworks. This guide demystifies AI for everyone—from software architects to policymakers and risk managers. Whether you're crafting LLMs or setting enterprise policy, this book empowers you to not only ask the right questions but build the right systems, because the future of AI isn't just about building intelligence, it's about building it responsibly.



153,50 €

153,50 € (zzgl. MwSt.)

vorbestellbar, Erscheinungstermin ca. März 2026

Artikelnummer: 9781003863397

Medium: Buch

ISBN: 978-1-003-86339-7

Verlag: Taylor & Francis Ltd

Erscheinungstermin: 05.03.2026

Sprache(n): Englisch

Auflage: 1. Auflage 2026

Produktform: Gebunden

Gewicht: 453 g

Seiten: 236

Format (B x H): 156 x 234 mm

