

ARTIFICIAL INTELLIGENCE

BENEFITS AND RISKS



*How machines learn, decide,
and change business*

**For professionals who make decisions
about AI — not build it**

DARKO MARTIĆ

CONTENTS

Preface

PART I WHAT IS ARTIFICIAL INTELLIGENCE?

- 1 Umbrellas, Wristwatches, and the Fear of New Machines
- 2 A Brief History of Thinking Machines
- 3 Can Machines Think? The Turing Test and the Chinese Room

PART II THE AI TOOLBOX

- 4 How Machines Learn: The Basics of Machine Learning
- 5 Deep Learning and Neural Networks
- 6 Teaching Computers to Read
- 7 Expert Systems: Bottled Knowledge
- 8 Evolution in a Computer: Genetic Algorithms
- 9 Computer Vision: Machines That See
- 10 Generative AI and Large Language Models
- 11 Sorting Out the Terms: AI, ML, Data Science, Deep Learning
- 12 The Hidden Enemy: Bias in Data

PART III MACHINE LEARNING IN PRACTICE

- 13 Supervised Learning: Learning from Examples
- 14 Unsupervised Learning: Finding Hidden Structure
- 15 Preparing the Data
- 16 Visual Data Analysis: The Orange Toolbox
- 17 Text Mining: From Raw Text to Insight

CONTENTS · CONTINUED

PART IV ALGORITHMS UP CLOSE

- 18 Linear and Logistic Regression
- 19 Finding the Groups: K-Means and Clustering
- 20 Decision Trees: Learning You Can Read
- 21 The Wisdom of the Crowd: Random Forests
- 22 Deep Learning in Action
- 23 How Good Is Your Model? Evaluation

PART V AI IN BUSINESS AND SOCIETY

- 24 The Business Benefits of AI
- 25 Benefits and Risks, Sector by Sector
- 26 The Ethics of Algorithms
- 27 The Algorithmic Society

PART VI THE PRACTITIONER'S COMPANION

- 28 The Churn Project: A Worked Example
- 29 Working with Generative AI: A Practical Playbook
- 30 Bringing AI into Your Organization
- 31 Deciding with Machines
- 32 The Road Ahead

CONTENTS · CONTINUED

PART VII FURTHER HORIZONS

- 33 Learning by Doing: Reinforcement Learning
- 34 Predicting Time: Forecasting and Time Series
- 35 When AI Is Attacked: Security Risks
- 36 Counting the Value: The Economics of AI Projects
- 37 Myths and Misconceptions
- 38 Auditing the Machines: Can We Verify AI?

Conclusion — Living and Working with AI

Appendix A — Questions for Reflection and Discussion

Further Reading

Appendix B — Hands-On with Orange

Appendix C — The Same Ideas in Excel

Glossary

Index

Acknowledgements

About the Author

PREFACE

Every generation meets a technology that seems to change the rules overnight. I was born in 1980. For my grandparents, it was electricity or the automobile. For my parents, the personal computer and the internet. For me, it is artificial intelligence — a family of technologies that can read, write, see, listen, predict, and more and more, decide.

And I need to know how it works.

PREFACE

Who This Book Is For

I wrote this book for people who are not computer scientists, but managers, analysts, teachers, students, and curious professionals who sense that AI matters to their work and their lives, but who have never been given a clear, jargon-free explanation of what it actually is and how it actually works.

PART I

What Is Artificial Intelligence?

New York, May 11, 1997. Garry Kasparov — by many measures the strongest chess player who ever lived — resigned the final game against IBM's Deep Blue, and for the first time a machine had defeated a reigning world champion in a match. Kasparov later said he had sensed “a new kind of intelligence” across the board. He was half right: there was no understanding in Deep Blue, only search and evaluation at inhuman speed. But the world's intuition shifted that day — from “machines cannot do that” to “what else can they do?” This part of the book is about that question: where it came from, and what it actually means.

WHAT YOU'LL LEARN

- Why fear of new technology is old, and usually aimed at the wrong thing
- Where AI actually came from — Turing, Dartmouth, and the AI winters
- The Turing test, the Chinese Room, and what “understanding” really means

Chapters 1-3

PART II

The AI Toolbox

Seoul, March 2016, game two of the match between Go legend Lee Sedol and the program AlphaGo. On move 37, the machine placed a stone so contrary to a thousand years of Go wisdom that commentators assumed it was a mistake, and Lee left the room to compose himself. It was not a mistake. Professionals now study “Move 37” as a masterpiece the human tradition had simply never explored. The machine had not been told the move; it had learned it, from millions of games played against itself. This part opens the toolbox that makes such learning possible, from humble regression to the neural networks behind that stone.

WHAT YOU'LL LEARN

- How machines learn from data, and what deep learning adds on top
- The toolbox: language, expert systems, genetic algorithms, computer vision, generative AI
- How to sort the jargon — AI, ML, deep learning, data science — and spot hidden bias

Chapters 4-12

PART III

Machine Learning in Practice

In a story made famous by a 2012 New York Times article, an American father walked into a Target store near Minneapolis, angry that the retailer was mailing his teenage daughter coupons for cribs and baby clothes. Days later, he called back to apologize himself: his daughter really was pregnant. Target's statisticians had learned that changes in purchases of a couple dozen ordinary products predicted pregnancy well enough to time the mailings. No one had told the model anything about anyone; the pattern was simply there, waiting. This part is about that craft: preparing data, and finding what it quietly contains.

WHAT YOU'LL LEARN

- Supervised vs. unsupervised learning, and when to use each
- How raw data becomes model-ready — cleaning, exploring, preparing
- Hands-on tools for visual data analysis and text mining

Chapters 13-17

PART IV

Algorithms Up Close

In 2006, Netflix offered one million dollars to any team that could improve its movie-recommendation model by ten percent. Three years and thousands of submissions later, a coalition of teams claimed the prize with an ensemble of over a hundred blended models. Then came the twist: Netflix never fully deployed the winning system — the engineering cost outweighed the value of its extra accuracy. The prize taught the industry two lessons at once: how far clever algorithms can push a metric, and that the best model on the leaderboard isn't always the best model for the business.

WHAT YOU'LL LEARN

- The core algorithms: regression, clustering, decision trees, random forests, deep learning
- How to judge whether a model actually deserves your trust
- A map for choosing the right tool for the problem in front of you

Chapters 18-23

PART V

AI in Business and Society

For several years in the mid-2010s, Amazon's engineers built an experimental system to rank job applicants' CVs, trained on a decade of the company's own hiring decisions. The system taught itself that male candidates were preferable — penalizing CVs that mentioned the word “women's,” and downgrading graduates of women's colleges. The engineers scrapped the project: nothing in the algorithm was malicious, it had simply distilled the history it was given, faithfully. Every theme of this part — benefit, risk, and responsibility — is contained in that one cancelled experiment.

WHAT YOU'LL LEARN

- The real business benefits AI delivers, and where they show up
- Benefit and risk sector by sector: finance, healthcare, education, employment, justice, media
- The ethics of algorithmic decisions, and what accountable AI looks like

Chapters 24-27

PART VI

The Practitioner's Companion

On the morning of August 1, 2012, the trading firm Knight Capital deployed new software to its automated trading systems. One of eight servers received the update incorrectly, and old code sprang to life, firing erroneous orders into the market at machine speed. In forty-five minutes, the firm lost approximately 440 million dollars and had to be rescued within the week. No model was malicious; no attacker was involved. Automation had simply outrun oversight. This part is the antidote: the practical craft of building, deploying, and governing systems that decide at machine speed — with humans still in charge.

WHAT YOU'LL LEARN

- A real project, start to finish: framing, building, deploying
- A practical playbook for working with generative AI day to day
- How to bring AI into an organization and keep humans in charge

Chapters 28–32

PART VII

Further Horizons

In 2019, the managing director of a British energy firm received a phone call from his boss at the German parent company, instructing him to urgently transfer €220,000 to a Hungarian supplier. He complied. The voice was a machine: criminals had cloned the executive's speech with AI and made off with the money. This final part explores the territory beyond the core curriculum: reinforcement learning, forecasting, the dark side of AI security, the economics of AI projects, and the myths that cloud public understanding.

WHAT YOU'LL LEARN

- Reinforcement learning and forecasting — learning from consequences, predicting the future
- The security risks of the AI era, and the economics behind every AI project
- The myths worth retiring, and how to audit an AI system you don't fully trust

Chapters 33-38

ACKNOWLEDGEMENTS

I did not start this book already knowing everything in it — I taught myself, testing tools and rewriting sections until I could explain them simply. Thank you to the researchers who publish openly, the makers of free tools such as Orange, and the people close to me, patient through many distracted evenings.

— *Darko Martić*

Machines now read scans, approve loans, rank job applications, and write fluent paragraphs. Most of us use them daily — without ever being told, clearly and honestly, how they work and when to trust them.

This book is the missing explanation: step by step, in plain language, from the first thinking machines to generative AI — with the practical judgment the algorithms cannot supply.

INSIDE, YOU WILL LEARN TO:

- ask the six questions that match any problem to the right tool
- read a model's confusion matrix the way a manager reads a P&L
- spot the hidden biases that turn honest data into unfair decisions
- use generative AI as a multiplier — without inheriting its mistakes
- bring AI into an organization and keep humans in charge of it

WHO THIS BOOK IS FOR:

- managers, analysts, and team leads who evaluate, buy, and deploy AI
- instructors and students of applied AI and business courses
- curious professionals who want substance — without code or heavy math

No programming, no heavy math — just the understanding professionals now need.



**MORE BOOKS BY
THE AUTHOR**

DARKO MARTIĆ · Talking in binary since 1980.

🌐 darko.martic.net

