

Machine Learning for Engineers


[Osvaldo Simeone](#), *King's College London*

Publication date

Digital publication date: 25 January 2023

Physical publication date: 03 November 2022

This self-contained introduction to machine learning, designed from the start with engineers in mind, will equip students with everything they need to start applying machine learning principles and algorithms to real-world engineering problems. With a consistent emphasis on the connections between estimation, detection, information theory, and optimization, it includes: an accessible overview of the relationships between machine learning and signal processing, providing a solid foundation for further study; clear explanations of the differences between state-of-the-art techniques and more classical methods, equipping students with all the understanding they need to make informed technique choices; demonstration of the links between information-theoretical concepts and their practical engineering relevance; reproducible examples using Matlab, enabling hands-on student experimentation. Assuming only a basic understanding of probability and linear algebra, and accompanied by lecture slides and solutions for instructors, this is the ideal introduction to machine learning for engineering students of all disciplines.



[Find out more](#)

Key features

- A book on machine learning written for engineers, by an engineer
- An accessible text with a unified information-theoretic framework
- Highlights connections between machine learning and estimation, detection, information theory, and optimization
- Offers concise but extensive coverage of state-of-the-art topics with simple, reproducible examples
- Derives modern methods, such as generative adversarial networks, from first principles, revealing their connection with standard techniques
- Divided into useful parts, allowing the book easily to be mapped to either a one- or a two-semester course

Resources

There are free resources and Instructor restricted resources available for this textbook [from the resources page](#).

About the book

Subjects: [Communications and Signal Processing](#), [Computer Science](#), [Engineering](#), [Machine Learning](#) and [Pattern Recognition](#)

Format: Hardback , Price: £54.99

Publication date: 03 November 2022, ISBN: 9781316512821

Format: Digital , Price: £54.99

Publication date: 25 January 2023, ISBN: 9781009072205

<https://doi.org/10.1017/9781009072205>

Contents

Part I - Introduction and Background

- 1 - When and How to Use Machine Learning
- 2 - Background

Part II - Fundamental Concepts and Algorithms

- 3 - Inference, or Model-Driven Prediction
- 4 - Supervised Learning: Getting Started
- 5 - Optimization for Machine Learning
- 6 - Supervised Learning: Beyond Least Squares
- 7 - Unsupervised Learning

Part III - Advanced Tools and Algorithms

- 8 - Statistical Learning Theory

9 - Exponential Family of Distributions

- 10 - Variational Inference and Variational Expectation Maximization

11 - Information-Theoretic Inference and Learning

- 12 - Bayesian Learning

Part IV - Beyond Centralized Single-Task Learning

- 13 - Transfer Learning, Multi-task Learning, Continual Learning, and Meta-learning

- 14 - Federated Learning

Part V - Epilogue

- 15 - Beyond This Book

Index