

Federated Learning

From Theory to Practice

How can we train powerful machine learning models together—across smartphones, hospitals, or financial institutions—without ever sharing raw data? This book delivers a compelling answer through the lens of federated learning (FL), a cutting-edge paradigm for decentralized, privacy-preserving machine learning. Designed for students, engineers, and researchers, this book offers a principled yet practical roadmap to building secure, scalable, and trustworthy FL systems from scratch.


At the heart of this book is a unifying framework that treats FL as a network-regularized optimization problem. This elegant formulation allows readers to seamlessly address personalization, robustness, and fairness—challenges often tackled in isolation. You'll learn how to structure FL networks based on task similarity, leverage graph-based methods and apply distributed optimization techniques to implement FL systems. Detailed pseudocode, intuitive explanations, and implementation-ready algorithms ensure you not only understand the theory but can apply it in real-world systems.

Topics such as privacy leakage analysis, model heterogeneity, and adversarial resilience are treated with both mathematical rigor and accessibility. Whether you're building decentralized AI for regulated industries or in settings where data, users, or system conditions change over time, this book equips you to design FL systems that are both performant and trustworthy.

How can we train powerful machine learning models together—across smartphones, hospitals, or financial institutions—without ever sharing raw data? This book delivers a compelling answer through the lens of federated learning (FL), a cutting-edge paradigm for decentralized, privacy-preserving machine learning. Designed for students, engineers, and researchers, this book offers a principled yet practical roadmap to building secure, scalable, and trustworthy FL systems from scratch.

At the heart of this book is a unifying framework that treats FL as a network-regularized optimization problem. This elegant formulation allows readers to seamlessly address personalization, robustness, and fairness—challenges often tackled in isolation. You'll learn how to structure FL networks based on task similarity, leverage graph-based methods and apply distributed optimization techniques to implement FL systems. Detailed pseudocode, intuitive explanations, and implementation-ready algorithms ensure you not only understand the theory but can apply it in real-world systems.

Topics such as privacy leakage analysis, model heterogeneity, and adversarial resilience are treated with both mathematical rigor and accessibility. Whether you're building decentralized AI for regulated industries or in settings where data, users, or system conditions change over time, this book equips you to design FL systems that are both performant and trustworthy.



69,54 €

64,99 € (zzgl. MwSt.)

vorbestellbar, Erscheinungstermin ca. Januar 2026

Artikelnummer: 9789819510085

Medium: Buch

ISBN: 978-981-951008-5

Verlag: Springer-Verlag GmbH

Erscheinungstermin: 03.01.2026

Sprache(n): Englisch

Auflage: Erscheinungsjahr 2026

Produktform: Gebunden

Gewicht: 516 g

Seiten: 213

Format (B x H): 160 x 241 mm

