Guerraoui / Gupta / Pinot

Robust Machine Learning

Distributed Methods for Safe AI

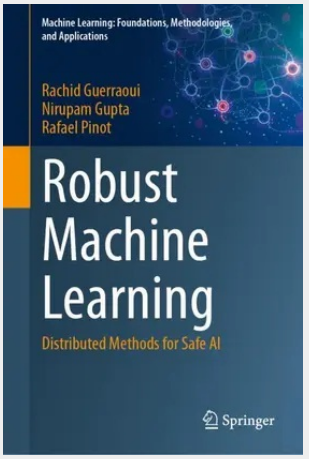

Today, machine learning algorithms are often distributed across multiple machines to leverage more computing power and more data. However, the use of a distributed framework entails a variety of security threats. In particular, some of the machines may misbehave and jeopardize the learning procedure. This could, for example, result from hardware and software bugs, data poisoning or a malicious player controlling a subset of the machines. This book explains in simple terms what it means for a distributed machine learning scheme to be robust to these threats, and how to build provably robust machine learning algorithms.

Studying the robustness of machine learning algorithms is a necessity given the ubiquity of these algorithms in both the private and public sectors. Accordingly, over the past few years, we have witnessed a rapid growth in the number of articles published on the robustness of distributed machine learning algorithms. We believe it is time to provide a clear foundation to this emerging and dynamic field. By gathering the existing knowledge and democratizing the concept of robustness, the book provides the basis for a new generation of reliable and safe machine learning schemes.

In addition to introducing the problem of robustness in modern machine learning algorithms, the book will equip readers with essential skills for designing distributed learning algorithms with enhanced robustness. Moreover, the book provides a foundation for future research in this area.

Today, machine learning algorithms are often distributed across multiple machines to leverage more computing power and more data. However, the use of a distributed framework entails a variety of security threats. In particular, some of the machines may misbehave and jeopardize the learning procedure. This could, for example, result from hardware and software bugs, data poisoning or a malicious player controlling a subset of the machines. This book explains in simple terms what it means for a distributed machine learning scheme to be robust to these threats, and how to build provably robust machine learning algorithms.Studying the robustness of machine learning algorithms is a necessity given the ubiquity of these algorithms in both the private and public sectors. Accordingly, over the past few years, we have witnessed a rapid growth in the number of articles published on the robustness of distributed machine learning algorithms. We believe it is time to provide a clear foundation to this emerging and dynamic field. By gathering the existing knowledge and democratizing the concept of robustness, the book provides the basis for a new generation of reliable and safe machine learning schemes. In addition to introducing the problem of robustness in modern machine learning algorithms, the book will equip readers with essential skills for designing distributed learning algorithms with enhanced robustness. Moreover, the book provides a foundation for future research in this area.

149,79 €
139,99 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9789819706877
Medium: Buch
ISBN: 978-981-97-0687-7
Verlag: Springer
Erscheinungstermin: 05.04.2024
Sprache(n): Englisch
Auflage: 2024
Serie: Machine Learning: Foundations, Methodologies, and Applications
Produktform: Gebunden
Gewicht: 453 g
Seiten: 170
Format (B x H): 160 x 241 mm

Kundenservice Verlag Dr. Otto Schmidt KG
Neumannstraße 10, 40235 Düsseldorf | kundenservice@otto-schmidt.de | 0221-93738-997 (Inland)