



1st ed. 2021, VIII, 202 p. 119 illus., 103 illus. in color.

#### Printed book

Softcover

54,99 € | £49.99 | \$69.99

<sup>[1]</sup>58,84 € (D) | 60,49 € (A) | CHF

65,00

#### eBook

46,00 € | £39.99 | \$54.99

<sup>[2]</sup>46,00 € (D) | 46,00 € (A) | CHF

52,00

Available from your library or  
[springer.com/shop](https://springer.com/shop)

#### MyCopy <sup>[3]</sup>

Printed eBook for just

€ | \$ 24.99

[springer.com/mycopy](https://springer.com/mycopy)

Leonida Gianfagna, Antonio Di Cecco

# Explainable AI with Python

- Offers a high-level perspective that explains the basics of XAI and its impacts on business and society, as well as a useful guide for machine learning practitioners to understand the current techniques to achieve explainability for AIML systems
- Fills the gaps to acquire the basic knowledge both from a theoretical and a practical perspective (with examples and direct implementation) making the reader quickly capable of working with tools and code for explainable AI
- Explains methods for the intrinsic interpretable ML models and agnostic methods for the non-interpretable ones

This book provides a full presentation of the current concepts and available techniques to make "machine learning" systems more explainable. The approaches presented can be applied to almost all the current "machine learning" models: linear and logistic regression, deep learning neural networks, natural language processing and image recognition, among the others. Progress in Machine Learning is increasing the use of artificial agents to perform critical tasks previously handled by humans (healthcare, legal and finance, among others). While the principles that guide the design of these agents are understood, most of the current deep-learning models are "opaque" to human understanding. Explainable AI with Python fills the current gap in literature on this emerging topic by taking both a theoretical and a practical perspective, making the reader quickly capable of working with tools and code for Explainable AI. Beginning with examples of what Explainable AI (XAI) is and why it is needed in the field, the book details different approaches to XAI depending on specific context and need. Hands-on work on interpretable models with specific examples leveraging Python are then presented, showing how intrinsic interpretable models can be interpreted and how to produce "human understandable" explanations. Model-agnostic methods for XAI are shown to produce explanations without relying on ML models internals that are "opaque."

Order online at [springer.com](https://springer.com) / or for the Americas call (toll free) 1-800-SPRINGER / or email us at: [customerservice@springernature.com](mailto:customerservice@springernature.com). / For outside the Americas call +49 (0) 6221-345-4301 / or email us at: [customerservice@springernature.com](mailto:customerservice@springernature.com).

The first € price and the £ and \$ price are net prices, subject to local VAT. Prices indicated with [1] include VAT for books; the €(D) includes 7% for Germany, the €(A) includes 10% for Austria. Prices indicated with [2] include VAT for electronic products; 19% for Germany, 20% for Austria. All prices exclusive of carriage charges. Prices and other details are subject to change without notice. All errors and omissions excepted. [3] No discount for MyCopy.

