

1. Record Nr.	UNINA9910736012003321
Autore	Sorvisto Dayne
Titolo	MLOps Lifecycle Toolkit : A Software Engineering Roadmap for Designing, Deploying, and Scaling Stochastic Systems // by Dayne Sorvisto
Pubbl/distr/stampa	Berkeley, CA : , : Apress : , : Imprint : Apress, , 2023
ISBN	9781484296424 1484296427
Edizione	[1st ed. 2023.]
Descrizione fisica	1 online resource (285 pages)
Disciplina	006.31
Soggetti	Machine learning Artificial intelligence Python (Computer program language) C++ (Computer program language) Machine Learning Artificial Intelligence Python C++
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Chapter 1: Introduction to Machine Learning Engineering -- Chapter 2: Developing Stochastic Systems -- Chapter 3: Tools for Data Science Developers -- Chapter 4: Infrastructure for MLOps -- Chapter 5, Building Training Pipelines -- Chapter 6: Building Inference Pipelines -- Chapter 7: Deploying Stochastic Systems -- Chapter 8: Data Ethics -- Chapter 9: Case Studies By Industry.
Sommario/riassunto	This book is aimed at practitioners of data science, with consideration for bespoke problems, standards, and tech stacks between industries. It will guide you through the fundamentals of technical decision making, including planning, building, optimizing, packaging, and deploying end-to-end, reliable, and robust stochastic workflows using the language of data science. MLOps Lifecycle Toolkit walks you through the principles of software engineering, assuming no prior

experience. It addresses the perennial “why” of MLOps early, along with insight into the unique challenges of engineering stochastic systems. Next, you’ll discover resources to learn software craftsmanship, data-driven testing frameworks, and computer science. Additionally, you will see how to transition from Jupyter notebooks to code editors, and leverage infrastructure and cloud services to take control of the entire machine learning lifecycle. You’ll gain insight into the technical and architectural decisions you’re likely to encounter, as well as best practices for deploying accurate, extensible, scalable, and reliable models. Through hands-on labs, you will build your own MLOps “toolkit” that you can use to accelerate your own projects. In later chapters, author Dayne Sorvisto takes a thoughtful, bottom-up approach to machine learning engineering by considering the hard problems unique to industries such as high finance, energy, healthcare, and tech as case studies, along with the ethical and technical constraints that shape decision making. After reading this book, whether you are a data scientist, product manager, or industry decision maker, you will be equipped to deploy models to production, understand the nuances of MLOps in the domain language of your industry, and have the resources for continuous delivery and learning.

You will:

- Understand the principles of software engineering and MLOps
- Design an end-to-end machine learning system
- Balance technical decisions and architectural trade-offs
- Gain insight into the fundamental problems unique to each industry and how to solve them.

---