

---

# Building Machine Learning Systems With Python Willi Richert

Recognizing the pretentiousness ways to acquire this books **Building Machine Learning Systems With Python Willi Richert** is additionally useful. You have remained in right site to begin getting this info. get the Building Machine Learning Systems With Python Willi Richert associate that we give here and check out the link.

You could buy lead Building Machine Learning Systems With Python Willi Richert or get it as soon as feasible. You could speedily download this Building Machine Learning Systems With Python Willi Richert after getting deal. So, in imitation of you require the books swiftly, you can straight get it. Its therefore totally simple and for that reason fats, isnt it? You have to favor to in this melody



[Building Machine Learning Powered Applications](#) "O'Reilly Media, Inc."

Unleash the power of Python and its robust data science capabilities  
About This Book Unleash the power of Python 3 objects Learn to use  
powerful Python libraries for effective data processing and analysis  
Harness the power of Python to analyze data and create insightful  
predictive models Unlock deeper insights into machine learning with this  
vital guide to cutting-edge predictive analytics Who This Book Is For  
Entry-level analysts who want to enter in the data science world will find  
this course very useful to get themselves acquainted with Python's data  
science capabilities for doing real-world data analysis. What You Will

Learn Install and setup Python Implement objects in Python by creating  
classes and defining methods Get acquainted with NumPy to use it with  
arrays and array-oriented computing in data analysis Create effective  
visualizations for presenting your data using Matplotlib Process and  
analyze data using the time series capabilities of pandas Interact with  
different kind of database systems, such as file, disk format, Mongo, and  
Redis Apply data mining concepts to real-world problems Compute on  
big data, including real-time data from the Internet Explore how to use  
different machine learning models to ask different questions of your data  
In Detail The Python: Real-World Data Science course will take you on a  
journey to become an efficient data science practitioner by thoroughly  
understanding the key concepts of Python. This learning path is divided  
into four modules and each module are a mini course in their own right,  
and as you complete each one, you'll have gained key skills and be ready  
for the material in the next module. The course begins with getting your  
Python fundamentals nailed down. After getting familiar with Python core  
concepts, it's time that you dive into the field of data science. In the second  
module, you'll learn how to perform data analysis using Python in a  
practical and example-driven way. The third module will teach you how to

---

design and develop data mining applications using a variety of datasets, starting with basic classification and affinity analysis to more complex data types including text, images, and graphs. Machine learning and predictive analytics have become the most important approaches to uncover data gold mines. In the final module, we'll discuss the necessary details regarding machine learning concepts, offering intuitive yet informative explanations on how machine learning algorithms work, how to use them, and most importantly, how to avoid the common pitfalls. Style and approach This course includes all the resources that will help you jump into the data science field with Python and learn how to make sense of data. The aim is to create a smooth learning path that will teach you how to get started with powerful Python libraries and perform various data science techniques in depth.

Large Scale Machine Learning with Python "O'Reilly Media, Inc."

"This video, with the help of practical projects, highlights how TensorFlow can be used in different scenarios--this includes projects for training models, machine learning, deep learning, and working with various neural networks. Each project provides exciting and insightful exercises that will teach you how to use TensorFlow and show you how layers of data can be explored by working with tensors"--Resource description page.

**Building Machine Learning and Deep Learning Models on Google Cloud Platform** "O'Reilly Media, Inc."

Learn to build powerful machine learning models quickly and deploy large-scale predictive applications About This Book Design, engineer and deploy scalable machine learning solutions

with the power of Python Take command of Hadoop and Spark with Python for effective machine learning on a map reduce framework Build state-of-the-art models and develop personalized recommendations to perform machine learning at scale Who This Book Is For This book is for anyone who intends to work with large and complex data sets. Familiarity with basic Python and machine learning concepts is recommended. Working knowledge in statistics and computational mathematics would also be helpful. What You Will Learn Apply the most scalable machine learning algorithms Work with modern state-of-the-art large-scale machine learning techniques Increase predictive accuracy with deep learning and scalable data-handling techniques Improve your work by combining the MapReduce framework with Spark Build powerful ensembles at scale Use data streams to train linear and non-linear predictive models from extremely large datasets using a single machine In Detail Large Python machine learning projects involve new problems associated with specialized machine learning architectures and designs that many data scientists have yet to tackle. But finding algorithms and designing and building platforms that deal with large sets of data is a growing need. Data scientists have to manage and maintain increasingly complex data projects, and with the rise of big data comes an increasing

---

demand for computational and algorithmic efficiency. Large Scale Machine Learning with Python uncovers a new wave of machine learning algorithms that meet scalability demands together with a high predictive accuracy. Dive into scalable machine learning and the three forms of scalability. Speed up algorithms that can be used on a desktop computer with tips on parallelization and memory allocation. Get to grips with new algorithms that are specifically designed for large projects and can handle bigger files, and learn about machine learning in big data environments. We will also cover the most effective machine learning techniques on a map reduce framework in Hadoop and Spark in Python. Style and Approach This efficient and practical title is stuffed full of the techniques, tips and tools you need to ensure your large scale Python machine learning runs swiftly and seamlessly. Large-scale machine learning tackles a different issue to what is currently on the market. Those working with Hadoop clusters and in data intensive environments can now learn effective ways of building powerful machine learning models from prototype to production. This book is written in a style that programmers from other languages (R, Julia, Java, Matlab) can follow.

**Machine Learning in Production** Building Machine Learning Systems with Python

Get more from your data by creating practical machine learning systems with Python Key Features Develop your own Python-based machine learning system Discover how Python offers multiple algorithms for modern machine learning systems Explore key Python machine learning libraries to implement in your projects Book Description Machine learning allows systems to learn things without being explicitly programmed to do so. Python is one of the most popular languages used to develop machine learning applications, which take advantage of its extensive library support. This third edition of Building Machine Learning Systems with Python addresses recent developments in the field by covering the most-used datasets and libraries to help you build practical machine learning systems. Using machine learning to gain deeper insights from data is a key skill required by modern application developers and analysts alike. Python, being a dynamic language, allows for fast exploration and experimentation. This book shows you exactly how to find patterns in your raw data. You will start by brushing up on your Python machine learning knowledge and being introduced to libraries. You'll quickly get to grips with serious, real-world projects on datasets, using modeling and creating recommendation systems. With Building Machine Learning Systems with Python, you'll gain the tools and understanding required to build your own systems, all tailored to solve real-world data analysis problems. By the end of this book, you will be able to build machine learning systems using techniques and methodologies such as classification, sentiment analysis, computer vision, reinforcement learning, and neural networks. What you will learn Build a

---

classification system that can be applied to text, images, and sound  
Employ Amazon Web Services (AWS) to run analysis on the  
cloud Solve problems related to regression using scikit-learn and  
TensorFlow Recommend products to users based on their past  
purchases Understand different ways to apply deep neural  
networks on structured data Address recent developments in the  
field of computer vision and reinforcement learning Who this  
book is for Building Machine Learning Systems with Python is for  
data scientists, machine learning developers, and Python  
developers who want to learn how to build increasingly complex  
machine learning systems. You will use Python's machine learning  
capabilities to develop effective solutions. Prior knowledge of  
Python progr ...

[Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow](#)  
Packt Publishing Ltd

Summary Machine Learning in Action is unique book that blends the  
foundational theories of machine learning with the practical realities of  
building tools for everyday data analysis. You'll use the flexible Python  
programming language to build programs that implement algorithms for  
data classification, forecasting, recommendations, and higher-level features  
like summarization and simplification. About the Book A machine is said to  
learn when its performance improves with experience. Learning requires  
algorithms and programs that capture data and ferret out the interesting or  
useful patterns. Once the specialized domain of analysts and mathematicians,  
machine learning is becoming a skill needed by many. Machine Learning in  
Action is a clearly written tutorial for developers. It avoids academic  
language and takes you straight to the techniques you'll use in your day-to-  
day work. Many (Python) examples present the core algorithms of statistical  
data processing, data analysis, and data visualization in code you can reuse.  
You'll understand the concepts and how they fit in with tactical tasks like

classification, forecasting, recommendations, and higher-level features like  
summarization and simplification. Readers need no prior experience with  
machine learning or statistical processing. Familiarity with Python is helpful.  
Purchase of the print book comes with an offer of a free PDF, ePub, and  
Kindle eBook from Manning. Also available is all code from the book. What's  
Inside A no-nonsense introduction Examples showing common ML tasks  
Everyday data analysis Implementing classic algorithms like Apriori and  
AdaBoos Table of Contents PART 1 CLASSIFICATION Machine learning  
basics Classifying with k-Nearest Neighbors Splitting datasets one feature at a  
time: decision trees Classifying with probability theory: naïve Bayes Logistic  
regression Support vector machines Improving classification with the  
AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES  
WITH REGRESSION Predicting numeric values: regression Tree-based  
regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled  
items using k-means clustering Association analysis with the Apriori algorithm  
Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL  
TOOLS Using principal component analysis to simplify data Simplifying data  
with the singular value decomposition Big data and MapReduce  
[Machine Learning with Python for Everyone](#) Apress

Learn the skills necessary to design, build, and deploy applications  
powered by machine learning (ML). Through the course of this hands-  
on book, you ' ll build an example ML-driven application from initial  
idea to deployed product. Data scientists, software engineers, and  
product managers—including experienced practitioners and novices  
alike—will learn the tools, best practices, and challenges involved in  
building a real-world ML application step by step. Author Emmanuel  
Ameisen, an experienced data scientist who led an AI education  
program, demonstrates practical ML concepts using code snippets,  
illustrations, screenshots, and interviews with industry leaders. Part I  
teaches you how to plan an ML application and measure success. Part  
II explains how to build a working ML model. Part III demonstrates

---

ways to improve the model until it fulfills your original vision. Part IV covers deployment and monitoring strategies. This book will help you:

- Define your product goal and set up a machine learning problem
- Build your first end-to-end pipeline quickly and acquire an initial dataset
- Train and evaluate your ML models and address performance bottlenecks
- Deploy and monitor your models in a production environment

Deep Learning Illustrated "O'Reilly Media, Inc."

Foundational Hands-On Skills for Succeeding with Real Data Science Projects This pragmatic book introduces both machine learning and data science, bridging gaps between data scientist and engineer, and helping you bring these techniques into production. It helps ensure that your efforts actually solve your problem, and offers unique coverage of real-world optimization in production settings. – From the Foreword by Paul Dix, series editor Machine Learning in Production is a crash course in data science and machine learning for people who need to solve real-world problems in production environments. Written for technically competent “accidental data scientists” with more curiosity and ambition than formal training, this complete and rigorous introduction stresses practice, not theory. Building on agile principles, Andrew and Adam Kelleher show how to quickly deliver significant value in production, resisting overhyped tools and unnecessary complexity. Drawing on their extensive experience, they help you ask useful questions and then execute production projects from start to finish. The authors show just how much information you can glean with straightforward queries, aggregations, and visualizations, and they teach

indispensable error analysis methods to avoid costly mistakes. They turn to workhorse machine learning techniques such as linear regression, classification, clustering, and Bayesian inference, helping you choose the right algorithm for each production problem. Their concluding section on hardware, infrastructure, and distributed systems offers unique and invaluable guidance on optimization in production environments. Andrew and Adam always focus on what matters in production: solving the problems that offer the highest return on investment, using the simplest, lowest-risk approaches that work. Leverage agile principles to maximize development efficiency in production projects Learn from practical Python code examples and visualizations that bring essential algorithmic concepts to life Start with simple heuristics and improve them as your data pipeline matures Avoid bad conclusions by implementing foundational error analysis techniques Communicate your results with basic data visualization techniques Master basic machine learning techniques, starting with linear regression and random forests Perform classification and clustering on both vector and graph data Learn the basics of graphical models and Bayesian inference Understand correlation and causation in machine learning models Explore overfitting, model capacity, and other advanced machine learning techniques Make informed architectural decisions about storage, data transfer, computation, and communication Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

Simon and Schuster

Building Machine Learning Systems with PythonPackt Publishing Ltd

---

## Python: Real-World Data Science Cambridge University Press

Get more from your data by creating practical machine learning systems with Python Key Features Develop your own Python-based machine learning system Discover how Python offers multiple algorithms for modern machine learning systems Explore key Python machine learning libraries to implement in your projects Book Description Machine learning allows systems to learn things without being explicitly programmed to do so. Python is one of the most popular languages used to develop machine learning applications, which take advantage of its extensive library support. This third edition of Building Machine Learning Systems with Python addresses recent developments in the field by covering the most-used datasets and libraries to help you build practical machine learning systems. Using machine learning to gain deeper insights from data is a key skill required by modern application developers and analysts alike. Python, being a dynamic language, allows for fast exploration and experimentation. This book shows you exactly how to find patterns in your raw data. You will start by brushing up on your Python machine learning knowledge and being introduced to libraries. You'll quickly get to grips with serious, real-world projects on datasets, using modeling and creating recommendation systems. With Building Machine Learning Systems with Python, you ' ll gain the tools and understanding required to build your own systems, all tailored to solve real-world data analysis problems. By the end of this book, you will be able to build machine learning systems using techniques and methodologies such as classification, sentiment analysis, computer vision, reinforcement learning, and neural networks. What you will learn Build a classification system that can be applied to text, images, and sound Employ Amazon Web Services (AWS) to run analysis on the cloud Solve problems related to regression using scikit-learn and TensorFlow Recommend products to users based on their past purchases Understand different ways to apply deep neural networks on structured data Address recent developments in the field of computer vision and reinforcement learning Who this book is for Building Machine Learning Systems with Python is for data scientists, machine learning developers, and Python developers who want to learn how to build increasingly complex

machine learning systems. You will use Python's machine learning capabilities to develop effective solutions. Prior knowledge of Python programming is expected.

Building Machine Learning Systems with Python Packt Publishing Ltd Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You ' ll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

## Distributed Machine Learning Patterns Packt Publishing Ltd

This is a tutorial-driven and practical, but well-grounded book showcasing good Machine Learning practices. There will be an emphasis on using existing technologies instead of showing how to write your own implementations of algorithms. This book is a scenario-based, example-driven tutorial. By the end of the book you will have learnt critical aspects of Machine Learning Python projects and experienced the power of ML-based systems by actually working on them. This book primarily targets Python developers who want to learn about and build Machine Learning into their projects, or who want to pro.

## Practical Machine Learning with Python John Wiley & Sons

---

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." – Tim Urban, author of *Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning*

Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. *Deep Learning Illustrated* is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn – with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens – presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners. Explore new tools that make deep learning models easier to build, use, and improve. Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more. Walk through building interactive

deep learning applications, and move forward with your own artificial intelligence projects. Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

*Designing Machine Learning Systems with Python* Pearson Professional

This book describes deep learning systems: the algorithms, compilers, and processor components to efficiently train and deploy deep learning models for commercial applications. The exponential growth in computational power is slowing at a time when the amount of compute consumed by state-of-the-art deep learning (DL) workloads is rapidly growing. Model size, serving latency, and power constraints are a significant challenge in the deployment of DL models for many applications. Therefore, it is imperative to co-design algorithms, compilers, and hardware to accelerate advances in this field with holistic system-level and algorithm solutions that improve performance, power, and efficiency. Advancing DL systems generally involves three types of engineers: (1) data scientists that utilize and develop DL algorithms in partnership with domain experts, such as medical, economic, or climate scientists; (2) hardware designers that develop specialized hardware to accelerate the components in the DL models; and (3) performance and compiler engineers that optimize software to run more efficiently on a given hardware. Hardware engineers should be aware of the characteristics and components of production and academic models likely to be adopted by industry to guide design decisions impacting future hardware. Data scientists should be aware of deployment platform constraints when designing models. Performance engineers should support optimizations across diverse models, libraries, and hardware targets. The purpose of this book is to provide a solid understanding of (1) the design, training, and applications of DL

---

algorithms in industry; (2) the compiler techniques to map deep learning code to hardware targets; and (3) the critical hardware features that accelerate DL systems. This book aims to facilitate co-innovation for the advancement of DL systems. It is written for engineers working in one or more of these areas who seek to understand the entire system stack in order to better collaborate with engineers working in other parts of the system stack. The book details advancements and adoption of DL models in industry, explains the training and deployment process, describes the essential hardware architectural features needed for today's and future models, and details advances in DL compilers to efficiently execute algorithms across various hardware targets. Unique in this book is the holistic exposition of the entire DL system stack, the emphasis on commercial applications, and the practical techniques to design models and accelerate their performance. The author is fortunate to work with hardware, software, data scientist, and research teams across many high-technology companies with hyperscale data centers. These companies employ many of the examples and methods provided throughout the book.

Machine Learning in Python Morgan & Claypool Publishers  
Design efficient machine learning systems that give you more accurate results  
About This Book Gain an understanding of the machine learning design process  
Optimize machine learning systems for improved accuracy  
Understand common programming tools and techniques for machine learning  
Develop techniques and strategies for dealing with large amounts of data from a variety of sources  
Build models to solve unique tasks  
Who This Book Is For This book is for data scientists, scientists, or just the curious. To get the most out of this book, you will need to know some linear algebra and some Python, and have a basic

knowledge of machine learning concepts. What You Will Learn  
Gain an understanding of the machine learning design process  
Optimize the error function of your machine learning system  
Understand the common programming patterns used in machine learning  
Discover optimizing techniques that will help you get the most from your data  
Find out how to design models uniquely suited to your task  
In Detail Machine learning is one of the fastest growing trends in modern computing. It has applications in a wide range of fields, including economics, the natural sciences, web development, and business modeling. In order to harness the power of these systems, it is essential that the practitioner develops a solid understanding of the underlying design principles. There are many reasons why machine learning models may not give accurate results. By looking at these systems from a design perspective, we gain a deeper understanding of the underlying algorithms and the optimisational methods that are available. This book will give you a solid foundation in the machine learning design process, and enable you to build customised machine learning models to solve unique problems. You may already know about, or have worked with, some of the off-the-shelf machine learning models for solving common problems such as spam detection or movie classification, but to begin solving more complex problems, it is important to adapt these models to your own specific needs. This book will give you this understanding and more. Style and approach This easy-to-follow, step-by-step guide covers the most important machine learning models and techniques from a design perspective.

Machine Learning Engineering Simon and Schuster



---

Machine learning has become an integral part of many commercial applications and research projects, but this field is not exclusive to large companies with extensive research teams. If you use Python, even as a beginner, this book will teach you practical ways to build your own machine learning solutions. With all the data available today, machine learning applications are limited only by your imagination. You'll learn the steps necessary to create a successful machine-learning application with Python and the scikit-learn library. Authors Andreas Müller and Sarah Guido focus on the practical aspects of using machine learning algorithms, rather than the math behind them. Familiarity with the NumPy and matplotlib libraries will help you get even more from this book. With this book, you'll learn: Fundamental concepts and applications of machine learning Advantages and shortcomings of widely used machine learning algorithms How to represent data processed by machine learning, including which data aspects to focus on Advanced methods for model evaluation and parameter tuning The concept of pipelines for chaining models and encapsulating your workflow Methods for working with text data, including text-specific processing techniques Suggestions for improving your machine learning and data science skills

**Machine Learning Systems** John Wiley & Sons

Roughly inspired by the human brain, deep neural networks trained with large amounts of data can solve complex tasks with unprecedented accuracy. This practical book provides an end-to-end guide to TensorFlow, the leading open source software library that helps you build and train neural networks for computer vision, natural language processing (NLP), speech recognition, and general predictive analytics. Authors Tom Hope, Yehezkel Resheff, and Itay Lieder provide a hands-on approach to TensorFlow fundamentals for a broad technical audience—from data scientists and engineers to students and

researchers. You'll begin by working through some basic examples in TensorFlow before diving deeper into topics such as neural network architectures, TensorBoard visualization, TensorFlow abstraction libraries, and multithreaded input pipelines. Once you finish this book, you'll know how to build and deploy production-ready deep learning systems in TensorFlow. Get up and running with TensorFlow, rapidly and painlessly Learn how to use TensorFlow to build deep learning models from the ground up Train popular deep learning models for computer vision and NLP Use extensive abstraction libraries to make development easier and faster Learn how to scale TensorFlow, and use clusters to distribute model training Deploy TensorFlow in a production setting

**Machine Learning in Action** Packt Publishing Ltd

Automate data and model pipelines for faster machine learning applications Key Features Build automated modules for different machine learning components Understand each component of a machine learning pipeline in depth Learn to use different open source AutoML and feature engineering platforms Book Description AutoML is designed to automate parts of Machine Learning. Readily available AutoML tools are making data science practitioners' work easy and are received well in the advanced analytics community. Automated Machine Learning covers the necessary foundation needed to create automated machine learning modules and helps you get up to speed with them in the most practical way possible. In this book, you'll learn how to automate different tasks in the machine learning pipeline such as data preprocessing, feature selection, model training, model optimization, and much more. In addition to this, it demonstrates how you can use the available automation

---

libraries, such as auto-sklearn and MLBox, and create and extend your own custom AutoML components for Machine Learning. By the end of this book, you will have a clearer understanding of the different aspects of automated Machine Learning, and you'll be able to incorporate automation tasks using practical datasets. You can leverage your learning from this book to implement Machine Learning in your projects and get a step closer to winning various machine learning competitions. What you will learn

- Understand the fundamentals of Automated Machine Learning systems
- Explore auto-sklearn and MLBox for AutoML tasks
- Automate your preprocessing methods along with feature transformation
- Enhance feature selection and generation using the Python stack
- Assemble individual components of ML into a complete AutoML framework
- Demystify hyperparameter tuning to optimize your ML models
- Dive into Machine Learning concepts such as neural networks and autoencoders
- Understand the information costs and trade-offs associated with AutoML

Who this book is for If you're a budding data scientist, data analyst, or Machine Learning enthusiast and are new to the concept of automated machine learning, this book is ideal for you. You'll also find this book useful if you're an ML engineer or data professional interested in developing quick machine learning pipelines for your projects. Prior exposure to Python programming will help you get the best out of this book.

Up and Running Google AutoML and AI Platform: Building Machine Learning and NLP Models Using AutoML and AI Platform for Production Environment (English Edition) Addison-Wesley Professional

Companies are spending billions on machine learning projects, but it's money wasted if the models can't be deployed effectively. In this practical

guide, Hannes Hapke and Catherine Nelson walk you through the steps of automating a machine learning pipeline using the TensorFlow ecosystem. You'll learn the techniques and tools that will cut deployment time from days to minutes, so that you can focus on developing new models rather than maintaining legacy systems. Data scientists, machine learning engineers, and DevOps engineers will discover how to go beyond model development to successfully productize their data science projects, while managers will better understand the role they play in helping to accelerate these projects.

- Understand the steps to build a machine learning pipeline
- Build your pipeline using components from TensorFlow Extended
- Orchestrate your machine learning pipeline with Apache Beam, Apache Airflow, and Kubeflow
- Pipelines Work with data using TensorFlow Data Validation and TensorFlow Transform
- Analyze a model in detail using TensorFlow Model Analysis
- Examine fairness and bias in your model performance
- Deploy models with TensorFlow Serving or TensorFlow Lite for mobile devices
- Learn privacy-preserving machine learning techniques

Deep Learning with TensorFlow and Keras True Positive Incorporated

Produce a fully functioning Intelligent System that leverages machine learning and data from user interactions to improve over time and achieve success. This book teaches you how to build an Intelligent System from end to end and leverage machine learning in practice. You will understand how to apply your existing skills in software engineering, data science, machine learning, management, and program management to produce working systems. Building Intelligent Systems is based on more than a decade of experience building Internet-scale Intelligent Systems that have hundreds of millions of user interactions per day in some of the largest and most important software systems in the world. What You'll Learn

- Understand the concept of an Intelligent System: What it is good for, when you need one, and how to set it up for success
- Design an intelligent user experience:

---

Produce data to help make the Intelligent System better over time  
Implement an Intelligent System: Execute, manage, and measure  
Intelligent Systems in practice Create intelligence: Use different  
approaches, including machine learning Orchestrate an Intelligent  
System: Bring the parts together throughout its life cycle and achieve  
the impact you want Who This Book Is For Software engineers,  
machine learning practitioners, and technical managers who want to  
build effective intelligent systems

Deep Learning with TensorFlow 2 and Keras BPB Publications

Practical patterns for scaling machine learning from your laptop to a  
distributed cluster. Scaling up models from standalone devices to large  
distributed clusters is one of the biggest challenges faced by modern machine  
learning practitioners. Distributed Machine Learning Patterns teaches you  
how to scale machine learning models from your laptop to large distributed  
clusters. In Distributed Machine Learning Patterns, you ' ll learn how to  
apply established distributed systems patterns to machine learning projects,  
and explore new ML-specific patterns as well. Firmly rooted in the real  
world, this book demonstrates how to apply patterns using examples based in  
TensorFlow, Kubernetes, Kubeflow, and Argo Workflows. Real-world  
scenarios, hands-on projects, and clear, practical DevOps techniques let you  
easily launch, manage, and monitor cloud-native distributed machine  
learning pipelines. Purchase of the print book includes a free eBook in PDF,  
Kindle, and ePub formats from Manning Publications.