Book Description

Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how.

By using concrete examples, minimal theory, and two production-ready python framework-Scikit-Learn and TensorFlow-author Aurelien Geron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You will learn a range of techniques, starting with simple liner regression and programming to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started.

Key Words

1. **Explore the machine learning landscape, particularly neural nets**
2. **Use Scikit-Learn to track an example machine learning project end-to-end.**
3. **Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods**
4. **Use the TensorFlow Library to build and train neural nets**
5. **Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning**
6. **Learn techniques for training and scaling deep neural nets**
7. **Apply practical code examples without acquiring excessive machine learning theory or algorithm details**

Table of Contents

Preface		**Xiii**
Part I.	**The Fundamentals of Machine Learning**	
01	**The Machine Learning Landscape**	**03**
02	**End-to-End Machine Learning Project**	**33**
03	**Classification**	**79**
04	**Training Models**	**105**

05	Support Vector Machines	145
06	Decision Trees	167
07	Ensemble Learning and Random Forests	181
08	Dimensionality Reduction	205
Part II. Neural Networks and Deep Learning		
09	Up and Running with TensorFlow	229
10	Introduction to Artificial Neural Networks	253
11	Training Deep Neural Nets	275
12	Distributing TensorFlow Across Devices and Servers	315
13	Convolutional Neural Networks	355
14	Recurrent Neural Networks	381
15	Autoencoders	415
16	Reinforcement Learning	441
A.	Exercise Solutions	475
B.	Machine Learning Project Checklist	501
C.	SVM Dual Problem	507
D.	Autodiff	511
E.	Other Popular ANN Architectures	519
Index		529

About Author

Aurelien Geron is a machine learning consultant. A former Googler, he led the Youtube video classification team from 2013 to 2016. He was also a founder and CTO of wifirst from 2002 to 2012, a leading wireless ISP in France, and a founder and CTO of polyconsell in 2001, the firm that now manages the electric car sharing service Autolib.