1.	Record Nr.	UNINA9910349527903321
	Autore	Michelucci Umberto
	Titolo	Advanced Applied Deep Learning : Convolutional Neural Networks and Object Detection / / by Umberto Michelucci
	Pubbl/distr/stampa	Berkeley, CA : , : Apress : , : Imprint : Apress, , 2019
	ISBN	9781484249765 1484249763
	Edizione	[1st ed. 2019.]
	Descrizione fisica	1 online resource (XVIII, 285 p. 88 illus., 28 illus. in color.)
	Disciplina	006.3
	Soggetti	Artificial intelligence Python (Computer program language) Open source software Artificial Intelligence Python Open Source
	Lingua di pubblicazione	Inglese
	Formato	Materiale a stampa
	Livello bibliografico	Monografia
	Note generali	Includes index.
	Nota di bibliografia	Includes bibliographical references.
	Nota di contenuto	Chapter 1: Introduction and Development Environment Setup -- Chapter 2: TensorFlow: advanced topics -- Chapter 3: Fundamentals of Convolutional Neural Networks -- Chapter 4: Advanced CNNs and Transfer Learning -- Chapter 5: Cost functions and style transfer -- Chapter 6: Object classification - an introduction -- Chapter 7: Object localization - an implementation in Python -- Chapter 8: Histology Tissue Classification.
	Sommario/riassunto	Develop and optimize deep learning models with advanced architectures. This book teaches you the intricate details and subtleties of the algorithms that are at the core of convolutional neural networks. In Advanced Applied Deep Learning, you will study advanced topics on CNN and object detection using Keras and TensorFlow. Along the way, you will look at the fundamental operations in CNN, such as convolution and pooling, and then look at more advanced architectures such as inception networks, resnets, and many more. While the book discusses theoretical topics, you will discover how to work efficiently with Keras with many tricks and tips, including how to customize

logging in Keras with custom callback classes, what is eager execution, and how to use it in your models. Finally, you will study how object detection works, and build a complete implementation of the YOLO (you only look once) algorithm in Keras and TensorFlow. By the end of the book you will have implemented various models in Keras and learned many advanced tricks that will bring your skills to the next level. You will: See how convolutional neural networks and object detection work Save weights and models on disk Pause training and restart it at a later stage Use hardware acceleration (GPUs) in your code Work with the Dataset TensorFlow abstraction and use pre-trained models and transfer learning Remove and add layers to pre-trained networks to adapt them to your specific project Apply pre-trained models such as Alexnet and VGG16 to new datasets.