Deep learning for computer vision with python kickstarter pdf

Continue

```
Instructor-led on a course scheduleOptional upgrade availableIf you have specific questions about this course, please contact us atsds-mm@mit.edu. Machine learning methods are commonly used across engineering and sciences, from computer systems (e.g.,
Netflix, Amazon), advertisers, and financial institutions employ machine learning algorithms for content recommendation, prediction or control. In this course, students
will learn about principles and algorithms for turning training data into effective automated predictions. We will cover: Representation, ver-fitting, regularization, probabilistic modeling, reinforcement learning; On-line algorithms, support vector machines, and neural
networks/deep learning. Students will implement and experiment with the algorithms in several Python projects designed for different practical applications. This course is part of the MITx MicroMasters Program in Statistics and Data Science. Master the skills needed to be an informed and effective practitioner of data science. You will complete this
course and three others from MITx, at a similar pace and level of rigor as an on-campus course at MIT, and then take a virtually-proctored exam to earn your path towards an MIT PhD or a Master's at other universities. To learn more about
this program, please visit. Please note: edX Inc. has recently entered into an agreement to transfer the edX platform to 2U, Inc., which will continue to run the platform thereafter. The sale will not affect your course enrollment, course fees or change your course fees or change your course experience for this offering. It is possible that the closing of the sale and the transfer of
the edX platform may be effectuated sometime in the Fall while this course is running. Please be aware that there could be changes to the edX platform Privacy of individual data for all learners who use the platform. For more information
see the edX Help Center. Understand principles behind machine learning problems such as classification, regression, clustering, and reinforcement learning Implement and analyze models for different applications Implement and organize machine
learning projects, from training, validation, parameter tuning, to feature engineering. Lectures: Introduction Linear classifiers, separability, perceptron algorithm Maximum margin hyperplane, loss, regularization Stochastic gradient descent, over-fitting, generalization Linear regression Recommender problems, collaborative filtering Non-linear
classification, kernels Learning features, Neural networks Deep learning, back propagation Recurrent neural networks Generalization, complexity, VC-dimension Unsupervised learning Reinforcement learning continued Applications:
Natural Language Processing Projects: Automatic Review Analyzer Digit Recognition with Neural Networks Reinforcement Learning Unfortunately, learners residing in one or more of the following countries or regions will not be able to register for this course: Iran, Cuba and the Crimea region of Ukraine. While edX has sought licenses from the U.S.
Office of Foreign Assets Control (OFAC) to offer our courses to learners in these countries and regions, the licenses we have received are not broad enough to allow us to offer this course to everyone, no matter where they live. 3rd Edition Released —
Officially supports Keras and TensorFlow 2.0. Grab your copy now! ...but you don't know how to get started. Let me help. $262,792 pledged of $5,000 goal I highly recommend grabbing a copy of Deep Learning for Computer Vision with Python. It goes into a lot of detail and has tons of detailed examples. It's the only book I've seen so far that covers
both how things work and how to actually use them in the real world to solve difficult problems. Check it out!" Adam GeitgeyAuthor of Machine Learning is Fun! blog series Phenomenal. The concepts on deep learning are so well explained that I will be recommending this book to anybody not just involved in computer vision but AI in general." Dr. Zig
ZdziarskiPhD in CV and ML, author at Zbigatron No problem, I have you covered either way. Regardless of your experience level, you'll find tremendous value inside Deep Learning for Computer Vision with Python, I guarantee it. Deep Learning for Computer Vision with Python will make you an expert in deep learning for computer vision and visual
recognition tasks. Inside the book we will focus on: Neural Networks and Machine Learning Convolutional Neural Networks (CNNs) Object detection, localization, and segmentation with deep learning Training large-scale (ImageNet-level) networks Hands on implementations using the Python programming language and the Keras, TensorFlow 2.0,
and mxnet deep learning libraries After going through Deep Learning for Computer Vision with Python, you'll be able to solve real-world problems with deep learning. I'm ready to order my copy now Python, TensorFlow 2.0, Keras, and mxnet are all well-built tools that, when combined, create a powerful deep learning development environment that
you can use to master deep learning for computer vision and visual recognition. Dr. Rosebrock delivers on what he promises! [Inside the book] he focuses on mastering deep learning and computer vision projects with detailed explanation of Python scripts, and puts to
your disposal priceless expertise to get you quickly engaged in the incredible field of deep learning. There is just no other book like this that I know of!" David BoulangerResearch Assistant in Data Analytics This book is for developers, researchers, and students who have at least some programming experience and want to become proficient in deep
learning for computer vision & visual recognition. If any of these descriptions fit you, rest assured: you're the target student. I wrote this book for you. You are a computer vision developer that utilizes OpenCV (among other image processing libraries) and are eager to level-up your skills. You have experience with machine learning and want to break
into neural networks/deep learning for image understanding. You are a scientist looking to apply deep learning + computer vision algorithms to your research. Your utilize computer vision algorithms in your own projects but have yet to try deep
learning. You used deep learning in projects before, but never in the context of visual recognition and image understanding. You write Python/machine learning hobbyist" who knows how to program and wants to understand what this "deep
learning" thing is all about. Download the FREE Deep Learning for Computer Vision Chapter! See for yourself how my book can help you master deep learning and computer vision. Adrian possesses a very rare talent of making complex concepts easy to grasp." Jean-Francois ParentPython Programmer at ChallengeU I'm constantly recommending
your [PyImageSearch.com] site to people I know at Georgia Tech and Udacity. While I consider Udacity the gold standard, I would rate your material at the same level. Keep up the good work." Andrew BakerSenior Software Engineer at Honeywell Since this book covers a huge amount of content, I've decided to break the book down into three
volumes called "bundles". A bundle includes all content from the lower volumes. You should choose a bundle based on: (1) how in-depth you want to study deep learning, computer vision, and visual recognition and (2) your particular
budget. Deep Learning for Computer Vision with Python is more than just a book. It's a complete package that is designed from the ground-up to help you master deep learning. Each bundle includes: The eBook files in PDF, .mobi, and .epub format. Video tutorials and walkthroughs for each chapter in the book. All source code listings so you can run
the examples from the book out-of-the-box. A downloadable pre-configured Ubuntu VirtualBox virtual machine that ships with all necessary Python + deep learning for Computer Vision with Python companion website, so you can further your knowledge, even when
you're done reading the book. The ImageNet Bundle also includes a physical hardcopy edition of the complete book delivered to your doorstep. Since this book covers a huge amount of content, I've decided to break the book down into three volumes called "bundles". Each bundle builds on top of the others and includes all content from the lower tiers.
Use the list of topics below (broken down by bundle) to help you (1) identify which topics you would like to study and then (2) choose a bundle based on this list. The Starter Bundle begins with a gentle introduction to the world of computer vision and machine learning, builds to neural networks, and then turns full steam into deep learning and
Convolutional Neural Networks. You'll even solve fun and interesting real-world problems using deep learning/neural networks. You are on a budget. While this is the lowest tier bundle, you'll still be getting a complete education. That said, for a more in-
depth treatment of deep learning for computer vision, I would recommend either the Practitioner Bundle or ImageNet Bundle. I'm ready to order my copy now Take Your First Steps Learn how to setup and configure your development environment to study deep learning using Python, TensorFlow 2.0, Keras, and mxnet. Understand Image Basics
Review how we represent images as arrays; coordinate systems; width, height, and depth; and aspect ratios. Machine Learning from data) and how we use data, feature vectors, scoring functions, and loss functions to create machine learning classifiers. Optimization Methods Gradient
Descent algorithms allow our algorithms to learn from data — I'll teach you how to implement then by hand. Feedforward Networks, which we'll discuss & implement the classic Perceptron algorithm, then move on to multi-layer networks, which we'll code from scratch via Python + Keras. Backpropagation Explained
We'll take an in-depth dive into the Backpropagation by hand using Python + NumPy. Intro to Convolutional Neural Networks (CNNs) I'll discuss exactly what a convolution is, followed by explaining Convolutional Neural Networks (what they are
used for, why they work so well for image classification, etc.). CNN Building Blocks Convolutional Neural Networks are built using different layers, dropout layers, activation layers, activation layers, batch normalization layers, to build your own CNNs. Uncover
Common Architectures & Training Patterns Discover common network architecture patterns you can use to design architectures of your own with minimal frustration and headaches. Pre-trained on 1,000 common object categories and are ready to be applied to
your own images/image datasets (VGG16, VGG19, ResNet50, etc.). Model Checkpointing Learn how to save and load your network models from disk during training, allowing you to checkpoint models from disk during training, allowing you to checkpoint models from disk during training.
determine if your network is underfitting or overfitting on your training data. Decay and Learning Rate Schedulers Learning rate decay/schedulers how to use these methods to maximize your model accuracy. Work With Your Own Datasets Learn how to gather your own
training images, label them, and train a Convolutional Neural Network from scratch on top of your dataset. LeNet Train the classic LeNet architecture from scratch to recognize handwritten digits in images. Case Study: Smile Detection I'll show you how to train a custom smile detector using Convolutional Neural Networks. The Practitioner Bundle is
appropriate if you want to take a deeper dive in deep learning. Inside this bundle, I cover more advanced techniques and best practices/rules of thumb. When you factor in the cost/time of training these deeper networks, the techniques and best practices/rules of thumb. When you factor in the cost/time of training these deeper networks, the techniques I cover in the Practitioner Bundle will save you so much time that the bundle will pay for itself, guaranteed. While
the Starter Bundle focuses on learning the fundamentals of deep learning, the Practitioner Bundle takes the next logical step and covers more advanced techniques, including transfer learning, fine-tuning, networks as feature extractors, working with HDF5 and large datasets, and object detection and localization. I also review Deep Dreaming and
Neural Style, Generative Adversarial Networks (GANs), and Image Super Resolution in detail. Using the techniques discussed in this bundle, you'll be able to compete in image classification competitions such as the Kaggle Dog vs. Cats Challenge (claiming a position in the top-25 leaderboard) and Stanford's cs231n Tiny ImageNet challenge. This
bundle is perfect for you if you are ready to study deep learning in-depth, understand advanced techniques, and discover common best practices and rules of thumb. I'm ready to order my copy now Everything in Starter Bundle Transfer Learning Don't train your CNN from scratch — use transfer learning and train your network in a fraction of the
time and obtain higher classification accuracy. Networks As Feature Extractors Treat pre-trained networks as feature extractors to obtain high classification accuracy with little effort. Fine-tuning to boost the accuracy of pre-trained networks, allowing you to work with small image dataset (and still reach high accuracy). Data
Augmentation Apply data augmentation to increase network classification accuracy without gathering more training data. Go Deeper Learn how to implement seminal CNN architectures from scratch, including AlexNet, VGGNet, SqueezeNet, and ResNet. Advanced Optimization Algorithms SGD is just the tip of the iceberg — you can also
train your networks using RMSprop, Adagrad, Adadelta, Adam, Adamax, and Nadam. I'll show you how. Over-sampling Utilize image cropping for an easy way to boost accuracy simply by training multiple networks. Best Practices
to Boost Network Performance Discover my optimal pathway for applying deep learning techniques to maximize classification accuracy (and which order to apply these techniques in to achieve the greatest effectiveness). Work With Datasets Too Large to Fit Into Memory Learn how to convert an image dataset from raw images on disk to HDF5
format, making networks easier (and faster) to train. Compete In Deep Learning Competitions I'll show you how to train a network on the Kaggle Dogs vs. Cats challenge and claim a position in the top-25 leaderboard with minimal effort. We'll also review how to rank high on the cs231n Tiny ImageNet classification challenge leaderboard. Object
Detection & Localization Discover how to use deep learning to detect and localize objects in images. Deep Dreaming and Neural Style Discover how to use deep learning to transform the artistic styles from one image to another. Generative Moversarial Networks (GANs) I'll show you how to utilize two neural networks (a generative model and a
discriminative model) to produce photorealistic images that look authentic to humans. Image Super Resolution input using deep learning algorithms. The Practitioner Bundle gives you the best bang for your buck. If you're even remotely serious about studying deep learning,
you should go with this bundle. The ImageNet Bundle is the most in-depth bundle and is a perfect fit if you want to train large-scale deep neural networks. This is also the only bundle that includes a hardcopy edition of the complete Deep Learning for Computer Vision with Python book, mailed to your doorstep. Inside this bundle, I demonstrate how to
build a custom Python framework to train networks. We'll use this framework to train AlexNet, VGGNet, SqueezeNet, GoogLeNet, and ResNet on the challenging ImageNet dataset. Using the training techniques I outline in this bundle, you'll be
able to reproduce the results you see in popular deep learning space. To demonstrate advanced deep learning techniques in action, I provide a number of case studies, including age + gender recognition, emotion and facial expression
recognition, car make + model recognition, and automatic image orientation (Mask R-CNN). Your copy of the ImageNet Bundle includes these bonus guides. You should choose the
ImageNet Bundle if: You want the complete deep learning for computer vision experience. Intend on training deep neural networks on large datasets from scratch. You want to learn how to train object detection or instance/semantic segmentation networks. I'm ready to order my copy now Everything in Starter Bundle Everything in Practitioner
Bundle ImageNet: Large Scale Visual Recognition Challenge to benchmark image classification algorithms. Work With ImageNet (1,000 categories) dataset is and why it's considered the de-facto challenge to benchmark image classification algorithms. Work With ImageNet I'll show you how to obtain the ImageNet dataset and convert it to an efficiently packed record file suitable
for training. Multi-GPU Training Learn how to utilize multiple GPUs to train your network in parallel, greatly reducing training time. AlexNet, VGGNet, GoogLeNet, SqueezeNet, and ResNet Train state-of-the-art network in parallel, greatly reducing training time. AlexNet, VGGNet, GoogLeNet, SqueezeNet, and ResNet Train state-of-the-art network in parallel, greatly reducing training time.
(and even months) of training time by discovering learning rate schedules that actually work — this chapter alone will save you enough time to actually pay for the book itself. Boost ImageNet Accuracy Learn how to restart training from saved epochs, lower learning rates, and increase classification accuracy on your testing set. Faster R-CNNs and
Single Shot Detectors (SSDs) In this bonus guide, I'll discuss object detection with deep learning, explain how the Faster R-CNN and Single Shot Detector (SSD) architectures work, and demonstrate how to use these architectures work, and demonstrate how to use these architectures using the Caffe framework. Case Study: Age + Gender Recognition Train your own custom CNN to (accurately) recognize
the age + gender of a person in an image using deep learning. Case Study: Emotion and Facial Expression Recognition Compete in Kaggle's Facial Expressions in real-time. Case Study: Vehicle Make + Model Classification Utilize fine-tuning to
train a network capable of recognizing the make + model of over 164 vehicles with over 96.52% accuracy. Case Study: Image Orientation but correct it as well. BONUS: Logo Detection with the RetinaNet Object
Detector Inside this chapter you will learn how to train the RetinaNet object detection framework to automatically detect logos in images with higher accuracy. BONUS: Weapon Detection Learn how to train an object detector capable of detection framework to automatically detect logos in images with higher accuracy. BONUS: Weapon Detection framework to automatically detect logos in images with higher accuracy.
you how to train a Mask R-CNN instance segmentation network to automatically detect skin lesions, a first step in cancer identification. I'll then show you how to train your own Mask R-CNN on your data. The core
deep learning guide Order My Copy Now The Starter Bundle is appropriate if (1) you are brand new to the world of computer vision and machine learning, builds to neural networks, and then turns full steam into deep learning and
Convolutional Neural Networks. Read More... While this is the lowest tier bundle, you'll still be getting a complete education. That said, for a more in-depth treatment of deep learning for computer vision, I would recommend either the Practitioner Bundle or ImageNet Bundle. All chapters inside the Starter Bundle in PDF, EPUB, and MOBI format.
Video tutorials and walkthroughs for each chapter. All source code listings so you can run the examples in the book out-of-the-box. A downloadable pre-configured Ubuntu VirtualBox virtual machine that ships with all necessary Python + deep learning for Computer
Vision with Python companion website. FREE updates as the book is revised and updated. Order My Copy Now or click here to pay with PayPal $295 Solve real-world problems using deep learning Order My Copy Now The Practitioner Bundle is geared towards readers who want an in-depth study of deep learning for computer vision. Here I cover
more advanced techniques and algorithms and demonstrate how to train networks to compete in popular image classification challenges. Read More... The Practitioner Bundle gives you the best bang for your buck. If you're even remotely serious about studying deep learning, you should go with this bundle. All chapters inside the Starter Bundle and
Practitioner Bundle in PDF, EPUB, and MOBI format. Video tutorials and walkthroughs for each chapter. All source code listings so you can run the examples in the book out-of-the-box. A downloadable pre-configured Ubuntu VirtualBox virtual machine that ships with all necessary Python + deep learning libraries you will need to be successful pre-
installed. Access to the Deep Learning for Computer Vision with Python companion website. FREE updates as the book is revised and updated. Order My Copy Now or click here to pay with PayPal $645 The most complete deep learning + computer vision education online today Order My Copy Now The ImageNet Bundle is the most in-depth bundle
and is for readers who want to train large-scale deep neural networks. This bundle is also the only bundle that includes a hardcopy edition of the complete Deep Learning for Computer Vision with Python framework to train network
architectures such as AlexNet, VGGNet, SqueezeNet, GoogLeNet, and ResNet from scratch on the challenging ImageNet dataset. You'll also learn how to train object detection (Faster R-CNNs, Single Shot Detectors, RetinaNet) and image segmentation (Mask R-CNN) networks on your own custom datasets. Using the training techniques I outline in
this bundle, you'll be able to reproduce the results you see in popular deep learning papers and publications — this is an absolute must for anyone doing research and development in the deep learning space. When it comes to studying deep learning, you can't beat this bundle! All chapters inside the Starter Bundle, Practitioner Bundle and ImageNet
Bundle in PDF, EPUB, and MOBI format. BONUS GUIDES on training object detection (Faster R-CNN, Single Shot Detector, RetinaNet) and image segmentation (Mask R-CNN) networks on your doorstep (this is the only bundle shot Detector).
that includes a physical copy of the book). Video tutorials and walkthroughs for each chapter. All source code listings so you can run the examples in the book out-of-the-box. A downloadable pre-configured Ubuntu VirtualBox virtual machine that ships with all necessary Python + deep learning libraries you will need to be successful pre-installed.
Access to the Deep Learning for Computer Vision with Python companion website. FREE updates as the book is revised and updated. Order My Copy Now or click here to pay with PayPal I can't recommend this book enough for anyone who has some basic knowledge of Python and is interested in Deep Learning and computer vision. Looking for anyone who has some basic knowledge of Python and is interested in Deep Learning and computer vision.
entry point to Deep Learning for image classification? Choose the Starter Bundle. Want to experiment with different well known architectures such as ResNet and GoogleNet? Go for the Practitioner Bundle. Want to train your networks on ImageNet? Go for the Practitioner Bundle. Want to experiment with different well known architectures such as ResNet and GoogleNet? Go for the Practitioner Bundle. Want to train your networks on ImageNet? Go for the Practitioner Bundle. Want to experiment with different well known architectures such as ResNet and GoogleNet? Go for the Practitioner Bundle. Want to experiment with different well known architectures such as ResNet and GoogleNet? Go for the Practitioner Bundle. Want to experiment with different well known architectures such as ResNet and GoogleNet? Go for the Practitioner Bundle. Want to experiment with different well known architectures such as ResNet and GoogleNet? Go for the Practitioner Bundle. Want to experiment with different well known architectures such as ResNet and GoogleNet? Go for the Practitioner Bundle. Want to experiment with different well known architectures such as ResNet and GoogleNet? Go for the Practitioner Bundle. Want to experiment with different well known architectures are the practical such as ResNet and GoogleNet? Go for the Practitioner Bundle. Want to experiment well known architectures are the practical such as ResNet and GoogleNet? Go for the Practitioner Bundle. Want to experiment well known architectures are the practical such as ResNet and GoogleNet?
Simon WalshConsultant Radiologist, Kings College Hospital I've learned a lot from the DL4CV book. I purchased an ImageNet bundle. I've learned a lot from the experiences you [Adrian] put in the book. Especially, I'm mind blown by the way you explained ResNet in Practitioner Bundle and Image Orientation correction from ImageNet bundle." Saideep
TalariComputer Vision Expert, Security Analyst at SenseHawk After reading my book, if you haven't learned the fundamentals of deep learning for computer vision, then I don't want your money. That's why I offer a 100% Money Back Guarantee. Simply send me an email and ask for a refund, up to 30 days after your purchase. With all the copies I've
sold, I count the number of refunds on one hand. My readers are satisfied and I'm sure you will be too. Your book is more expensive than other online courses and books — why is your book worth the price? First, it's important to understand that Deep Learning for Computer Vision with Python is the most complete, comprehensive deep learning
education online (the ImageNet Bundle is over 900+ pages). Not only does it cover the theory behind deep learning, it also details the implementation as well. You can't find a book this detailed in any other online platform. MOOC, or book. Secondly, I personally dedicate time daily to answering your guestions, providing help, and offering suggestions
— no other book or course online gives you this level of access to authors. To be totally honest with you, I've considered raising the price of this book multiple times but haven't (yet). My book may seem expensive, but the value you are getting is multiple orders of magnitude higher than any other book or course. I encourage you to give my book a try.
Once you dig into the content I'm confident you'll agree that the book is well worth the price. What deep learning researcher and practitioner, I can tell you that the combination of Keras and TensorFlow 2.0 is the
fastest, easiest way to go from idea, to experimentation, to result. The mxnet library specializes in distributed learning networks. What if I'm a
beginner at deep learning? Don't worry; you won't get bogged down by tons of theory and complex equations. We'll start off with the basics of machine learning and neural networks. You'll be a neural network ninja in no time, and be able to graduate to the more advanced content. Do I need any
programming experience before reading this book? This book assumes you have some prior programming experience (e.g. you know what a variable, function, loop, etc. are). You should have more skills than a novice, but certainly not an intermediate or advanced developer. As long as you understand basic programming logic flow you'll be successful
in reading (and understanding) the contents of this book. Do I need any special hardware to run the examples in the Starter Bundle or Practitioner Bundle, although some examples will take longer to
run. In either case, a GPU will dramatically speed up the network training process but is not a requirement. Can I upgrade from a lower tier bundle to a higher one. For example, you could purchase the Starter Bundle and then upgrade to the Practitioner
Bundle or ImageNet Bundle at a later date. The cost to upgrade would simply be the price difference between your current bundle and the bundle you wanted to upgrade would not need to "repurchase" the content you already own). To upgrade your bundle just send me an email and I can get you the upgrade link. What happens after I
purchase? After you purchase your copy of Deep Learning for Computer Vision with Python you will (1) receive an email receive an email to
enter your shipping information. Why are we using the Python programming language? First of all, Python is awesome. It is an easy language to learn and hands-down the best way to work with deep learning algorithms. The simple, intuitive syntax allows you to focus on learning the basics of deep learning, rather than spending hours fixing crazy
compiler errors in other languages. I want to learn TensorFlow 2.0 is covered? Yes, TensorFlow 2.0 is covered inside the text. We primarily use TensorFlow 2.0 and the Keras API inside TensorFlow 2.0 specific features such as
GradientTape and eager execution. What if I'm already experienced in deep learning? This book isn't just for beginners — there's advanced content in here too. You'll build a custom framework that can be used to train very deep architectures on the challenging
ImageNet dataset from scratch. I'll even show you my personal blueprint that I use to determine which deep learning techniques to apply when confronted with a new problem. Best of all, these solutions and tactics can be directly applied to your current job, research, and projects. Do I need to know OpenCV? You do not need to know the OpenCV
library to be successful when going through this book. We only use OpenCV to facilitate basic image processing operations. That said, a little bit of OpenCV experience goes a long way, so if you're new to OpenCV I highly recommend (1) purchase a copy
of Deep Learning for Computer Vision with Python and (2) work through my other book, Practical Python and OpenCV to learn the fundamentals. Do I need special hardware for the ImageNet Bundle, the
better. You should also have at least 1TB of free space on your machine. The ImageNet Bundle covers very advanced deep learning techniques on massive datasets, so make sure you make the necessary hardware preparations. What GPUs do you recommend for the book? I personally use the NVIDIA Titan X (12GB) on a daily basis for training my own
deep learning networks. The Titan X a bit expensive, so NVIDIA has released the GTX 1080 with 8GB of memory for half the cost of the Titan X. The latest addition to the NVIDIA family, the 1080 Ti (11GB), is also highly recommended. Alternatively, I would recommend using Amazon EC2 and their GPU instances (particularly p2.* and q2.*) in the
cloud to train your networks if you do not want to purchase physical hardware. Can I use the cloud for deep learning? Yes, you can absolutely use cloud services such as Amazon Web Services (AWS) or Microsoft Azure either with or without a GPU to work through the examples in this book. To jumpstart your education, I have released my own
personal pre-configured Amazon Machine Instance (AMI) to help you with your studies and projects. Simply launch an EC2 instance using this pre-configured AMI and you'll be ready to train your own deep neural networks in the matter of minutes! If you're a Microsoft Azure user, you can spin up Microsft DSVM instance and be up and running in a
few minutes as well. Are the hardcopy editions shipping? Yep, the hardcopy edition. After you purchase, you will receive an email with a link to enter your shipping information. Once I have your shipping address I can get your hardcopy edition in the mail,
normally within 48 hours. Where can I learn more about you? I have authored over 350+ blog posts about computer vision, OpenCV, and deep learning over at PyImageSearch.com. Check out the posts to get a feel for my teaching and writing style (not to mention the quality and depth of the tutorials). I would also highly suggest that you sign up for
the (free) Table of Contents and sample chapters I am offering using the form at the bottom-right corner of this page. Who's behind this? Hey, I'm Adrian Rosebrock, a Ph.D and entrepreneur who has spent his entire adult life studying computer vision and machine learning. Over the past 6 years alone I have: Started the PyImageSearch.com blog and
published over 350+ tutorials and articles aimed at teaching computer vision, deep learning, and OpenCV. This course is the most computer vision education online today, covering 13 modules broken out into 168 lessons with over 2,161
pages of content. Authored Deep Learning for Computer Vision with Python, the most in-depth computer vision and deep learning book available today, including super practical walkthroughs, hands-on tutorials (with lots of code), and a no-nonsense teaching style that will help you master computer vision and deep learning. Published Raspberry Pi
for Computer Vision, which covers embedded computer vision and deep learning on devices such as the Raspberry Pi, Google Coral, Movidius NCS, and NVIDIA Jetson Nano. Answered over 50,000+ emails and helped 10,000s of developers, researchers, and students just like yourself learn the ropes of computer vision and deep learning. If studying
deep learning and visual recognition sounds interesting to you, I hope you'll consider grabbing a copy of this book. You'll learn a ton about deep learning and computer vision in a practical, hands-on way. And you'll have fun doing it. See you on the other side!
```