

# Contents

<b>1</b>	<b>Introduction</b> .....	<b>15</b>
<b>1.1</b>	<b>I Studied Deep Learning the <i>Wrong Way</i>... This Is the <i>Right Way</i></b> .....	<b>15</b>
<b>1.2</b>	<b>Who This Book Is For</b> .....	<b>17</b>
1.2.1	Just Getting Started in Deep Learning? .....	17
1.2.2	Already a Seasoned Deep Learning Practitioner? .....	17
<b>1.3</b>	<b>Book Organization</b> .....	<b>17</b>
1.3.1	Volume #1: Starter Bundle .....	17
1.3.2	Volume #2: Practitioner Bundle .....	18
1.3.3	Volume #3: ImageNet Bundle .....	18
1.3.4	Need to Upgrade Your Bundle? .....	18
<b>1.4</b>	<b>Tools of the Trade: Python, Keras, and Mxnet</b> .....	<b>18</b>
1.4.1	What About TensorFlow? .....	18
1.4.2	Do I Need to Know OpenCV? .....	19
<b>1.5</b>	<b>Developing Our Own Deep Learning Toolset</b> .....	<b>19</b>
<b>1.6</b>	<b>Summary</b> .....	<b>20</b>
<b>2</b>	<b>What Is Deep Learning?</b> .....	<b>22</b>
<b>2.1</b>	<b>A Concise History of Neural Networks and Deep Learning</b> .....	<b>22</b>
<b>2.2</b>	<b>Hierarchical Feature Learning</b> .....	<b>24</b>
<b>2.3</b>	<b>How "Deep" Is Deep?</b> .....	<b>27</b>
<b>2.4</b>	<b>Summary</b> .....	<b>30</b>
<b>3</b>	<b>Image Fundamentals</b> .....	<b>31</b>
<b>3.1</b>	<b>Pixels: The Building Blocks of Images</b> .....	<b>31</b>
3.1.1	Forming an Image From Channels .....	34

<b>3.2</b>	<b>The Image Coordinate System</b>	<b>34</b>
3.2.1	Images as NumPy Arrays	35
3.2.2	RGB and BGR Ordering	36
<b>3.3</b>	<b>Scaling and Aspect Ratios</b>	<b>36</b>
<b>3.4</b>	<b>Summary</b>	<b>38</b>
<b>4</b>	<b>Image Classification Basics</b>	<b>39</b>
<b>4.1</b>	<b>What Is Image Classification?</b>	<b>40</b>
4.1.1	A Note on Terminology	40
4.1.2	The Semantic Gap	41
4.1.3	Challenges	42
<b>4.2</b>	<b>Types of Learning</b>	<b>45</b>
4.2.1	Supervised Learning	45
4.2.2	Unsupervised Learning	46
4.2.3	Semi-supervised Learning	47
<b>4.3</b>	<b>The Deep Learning Classification Pipeline</b>	<b>48</b>
4.3.1	A Shift in Mindset	48
4.3.2	Step #1: Gather Your Dataset	50
4.3.3	Step #2: Split Your Dataset	50
4.3.4	Step #3: Train Your Network	51
4.3.5	Step #4: Evaluate	51
4.3.6	Feature-based Learning versus Deep Learning for Image Classification	51
4.3.7	What Happens When my Predictions Are Incorrect?	52
<b>4.4</b>	<b>Summary</b>	<b>52</b>
<b>5</b>	<b>Datasets for Image Classification</b>	<b>53</b>
<b>5.1</b>	<b>MNIST</b>	<b>53</b>
<b>5.2</b>	<b>Animals: Dogs, Cats, and Pandas</b>	<b>54</b>
<b>5.3</b>	<b>CIFAR-10</b>	<b>55</b>
<b>5.4</b>	<b>SMILES</b>	<b>55</b>
<b>5.5</b>	<b>Kaggle: Dogs vs. Cats</b>	<b>56</b>
<b>5.6</b>	<b>Flowers-17</b>	<b>57</b>
<b>5.7</b>	<b>CALTECH-101</b>	<b>57</b>
<b>5.8</b>	<b>Tiny ImageNet 200</b>	<b>58</b>
<b>5.9</b>	<b>Adience</b>	<b>58</b>
<b>5.10</b>	<b>ImageNet</b>	<b>58</b>
5.10.1	What Is ImageNet?	59
5.10.2	ImageNet Large Scale Visual Recognition Challenge (ILSVRC)	59
<b>5.11</b>	<b>Kaggle: Facial Expression Recognition Challenge</b>	<b>59</b>
<b>5.12</b>	<b>Indoor CVPR</b>	<b>60</b>
<b>5.13</b>	<b>Stanford Cars</b>	<b>60</b>
<b>5.14</b>	<b>LISA Traffic Signs</b>	<b>60</b>
<b>5.15</b>	<b>Front/Rear View Vehicles</b>	<b>61</b>
<b>5.16</b>	<b>Summary</b>	<b>62</b>

<b>6</b>	<b>Configuring Your Development Environment</b>	<b>63</b>
<b>6.1</b>	<b>Libraries and Packages</b>	<b>63</b>
6.1.1	Python	63
6.1.2	Keras	64
6.1.3	Mxnet	64
6.1.4	OpenCV, scikit-image, scikit-learn, and more	64
<b>6.2</b>	<b>Configuring Your Development Environment?</b>	<b>64</b>
<b>6.3</b>	<b>Preconfigured Virtual Machine</b>	<b>65</b>
<b>6.4</b>	<b>Cloud-based Instances</b>	<b>65</b>
<b>6.5</b>	<b>How to Structure Your Projects</b>	<b>65</b>
<b>6.6</b>	<b>Summary</b>	<b>66</b>
<b>7</b>	<b>Your First Image Classifier</b>	<b>67</b>
<b>7.1</b>	<b>Working with Image Datasets</b>	<b>67</b>
7.1.1	Introducing the “Animals” Dataset	67
7.1.2	The Start to Our Deep Learning Toolkit	68
7.1.3	A Basic Image Preprocessor	69
7.1.4	Building an Image Loader	70
<b>7.2</b>	<b>k-NN: A Simple Classifier</b>	<b>72</b>
7.2.1	A Worked k-NN Example	74
7.2.2	k-NN Hyperparameters	75
7.2.3	Implementing k-NN	75
7.2.4	k-NN Results	78
7.2.5	Pros and Cons of k-NN	79
<b>7.3</b>	<b>Summary</b>	<b>80</b>
<b>8</b>	<b>Parameterized Learning</b>	<b>81</b>
<b>8.1</b>	<b>An Introduction to Linear Classification</b>	<b>82</b>
8.1.1	Four Components of Parameterized Learning	82
8.1.2	Linear Classification: From Images to Labels	83
8.1.3	Advantages of Parameterized Learning and Linear Classification	84
8.1.4	A Simple Linear Classifier With Python	85
<b>8.2</b>	<b>The Role of Loss Functions</b>	<b>88</b>
8.2.1	What Are Loss Functions?	88
8.2.2	Multi-class SVM Loss	89
8.2.3	Cross-entropy Loss and Softmax Classifiers	91
<b>8.3</b>	<b>Summary</b>	<b>94</b>
<b>9</b>	<b>Optimization Methods and Regularization</b>	<b>95</b>
<b>9.1</b>	<b>Gradient Descent</b>	<b>96</b>
9.1.1	The Loss Landscape and Optimization Surface	96
9.1.2	The “Gradient” in Gradient Descent	97
9.1.3	Treat It Like a Convex Problem (Even if It’s Not)	98
9.1.4	The Bias Trick	98
9.1.5	Pseudocode for Gradient Descent	99
9.1.6	Implementing Basic Gradient Descent in Python	100

9.1.7	Simple Gradient Descent Results	104
<b>9.2</b>	<b>Stochastic Gradient Descent (SGD)</b>	<b>106</b>
9.2.1	Mini-batch SGD	106
9.2.2	Implementing Mini-batch SGD	107
9.2.3	SGD Results	111
<b>9.3</b>	<b>Extensions to SGD</b>	<b>111</b>
9.3.1	Momentum	112
9.3.2	Nesterov's Acceleration	113
9.3.3	Anecdotal Recommendations	113
<b>9.4</b>	<b>Regularization</b>	<b>114</b>
9.4.1	What Is Regularization and Why Do We Need It?	114
9.4.2	Updating Our Loss and Weight Update To Include Regularization	115
9.4.3	Types of Regularization Techniques	117
9.4.4	Regularization Applied to Image Classification	117
<b>9.5</b>	<b>Summary</b>	<b>119</b>
<b>10</b>	<b>Neural Network Fundamentals</b>	<b>121</b>
<b>10.1</b>	<b>Neural Network Basics</b>	<b>121</b>
10.1.1	Introduction to Neural Networks	122
10.1.2	The Perceptron Algorithm	129
10.1.3	Backpropagation and Multi-layer Networks	137
10.1.4	Multi-layer Networks with Keras	153
10.1.5	The Four Ingredients in a Neural Network Recipe	163
10.1.6	Weight Initialization	165
10.1.7	Constant Initialization	165
10.1.8	Uniform and Normal Distributions	165
10.1.9	LeCun Uniform and Normal	166
10.1.10	Glorot/Xavier Uniform and Normal	166
10.1.11	He et al./Kaiming/MSRA Uniform and Normal	167
10.1.12	Differences in Initialization Implementation	167
<b>10.2</b>	<b>Summary</b>	<b>168</b>
<b>11</b>	<b>Convolutional Neural Networks</b>	<b>169</b>
<b>11.1</b>	<b>Understanding Convolutions</b>	<b>170</b>
11.1.1	Convolutions versus Cross-correlation	170
11.1.2	The "Big Matrix" and "Tiny Matrix" Analogy	171
11.1.3	Kernels	171
11.1.4	A Hand Computation Example of Convolution	172
11.1.5	Implementing Convolutions with Python	173
11.1.6	The Role of Convolutions in Deep Learning	179
<b>11.2</b>	<b>CNN Building Blocks</b>	<b>179</b>
11.2.1	Layer Types	181
11.2.2	Convolutional Layers	181
11.2.3	Activation Layers	186
11.2.4	Pooling Layers	186
11.2.5	Fully-connected Layers	188
11.2.6	Batch Normalization	189
11.2.7	Dropout	190

<b>11.3</b>	<b>Common Architectures and Training Patterns</b>	<b>191</b>
11.3.1	Layer Patterns	191
11.3.2	Rules of Thumb	192
<b>11.4</b>	<b>Are CNNs Invariant to Translation, Rotation, and Scaling?</b>	<b>194</b>
<b>11.5</b>	<b>Summary</b>	<b>195</b>
<b>12</b>	<b>Training Your First CNN</b>	<b>197</b>
<b>12.1</b>	<b>Keras Configurations and Converting Images to Arrays</b>	<b>197</b>
12.1.1	Understanding the keras.json Configuration File	197
12.1.2	The Image to Array Preprocessor	198
<b>12.2</b>	<b>ShallowNet</b>	<b>200</b>
12.2.1	Implementing ShallowNet	200
12.2.2	ShallowNet on Animals	202
12.2.3	ShallowNet on CIFAR-10	206
<b>12.3</b>	<b>Summary</b>	<b>209</b>
<b>13</b>	<b>Saving and Loading Your Models</b>	<b>211</b>
<b>13.1</b>	<b>Serializing a Model to Disk</b>	<b>211</b>
<b>13.2</b>	<b>Loading a Pre-trained Model from Disk</b>	<b>214</b>
<b>13.3</b>	<b>Summary</b>	<b>217</b>
<b>14</b>	<b>LeNet: Recognizing Handwritten Digits</b>	<b>219</b>
<b>14.1</b>	<b>The LeNet Architecture</b>	<b>219</b>
<b>14.2</b>	<b>Implementing LeNet</b>	<b>220</b>
<b>14.3</b>	<b>LeNet on MNIST</b>	<b>222</b>
<b>14.4</b>	<b>Summary</b>	<b>227</b>
<b>15</b>	<b>MiniVGGNet: Going Deeper with CNNs</b>	<b>229</b>
<b>15.1</b>	<b>The VGG Family of Networks</b>	<b>229</b>
15.1.1	The (Mini) VGGNet Architecture	230
<b>15.2</b>	<b>Implementing MiniVGGNet</b>	<b>230</b>
<b>15.3</b>	<b>MiniVGGNet on CIFAR-10</b>	<b>234</b>
15.3.1	With Batch Normalization	236
15.3.2	Without Batch Normalization	237
<b>15.4</b>	<b>Summary</b>	<b>238</b>
<b>16</b>	<b>Learning Rate Schedulers</b>	<b>241</b>
<b>16.1</b>	<b>Dropping Our Learning Rate</b>	<b>241</b>
16.1.1	The Standard Decay Schedule in Keras	242
16.1.2	Step-based Decay	244
16.1.3	Implementing Custom Learning Rate Schedules in Keras	245
<b>16.2</b>	<b>Summary</b>	<b>249</b>

<b>17</b>	<b>Spotting Underfitting and Overfitting</b>	<b>251</b>
<b>17.1</b>	<b>What Are Underfitting and Overfitting?</b>	<b>251</b>
17.1.1	Effects of Learning Rates	253
17.1.2	Pay Attention to Your Training Curves	254
17.1.3	What if Validation Loss Is Lower than Training Loss?	254
<b>17.2</b>	<b>Monitoring the Training Process</b>	<b>255</b>
17.2.1	Creating a Training Monitor	255
17.2.2	Babysitting Training	257
<b>17.3</b>	<b>Summary</b>	<b>260</b>
<b>18</b>	<b>Checkpointing Models</b>	<b>263</b>
<b>18.1</b>	<b>Checkpointing Neural Network Model Improvements</b>	<b>263</b>
<b>18.2</b>	<b>Checkpointing Best Neural Network Only</b>	<b>267</b>
<b>18.3</b>	<b>Summary</b>	<b>269</b>
<b>19</b>	<b>Visualizing Network Architectures</b>	<b>271</b>
<b>19.1</b>	<b>The Importance of Architecture Visualization</b>	<b>271</b>
19.1.1	Installing graphviz and pydot	272
19.1.2	Visualizing Keras Networks	272
<b>19.2</b>	<b>Summary</b>	<b>275</b>
<b>20</b>	<b>Out-of-the-box CNNs for Classification</b>	<b>277</b>
<b>20.1</b>	<b>State-of-the-art CNNs in Keras</b>	<b>277</b>
20.1.1	VGG16 and VGG19	278
20.1.2	ResNet	279
20.1.3	Inception V3	280
20.1.4	Xception	280
20.1.5	Can We Go Smaller?	280
<b>20.2</b>	<b>Classifying Images with Pre-trained ImageNet CNNs</b>	<b>281</b>
20.2.1	Classification Results	284
<b>20.3</b>	<b>Summary</b>	<b>286</b>
<b>21</b>	<b>Case Study: Breaking Captchas</b>	<b>287</b>
<b>21.1</b>	<b>Breaking Captchas with a CNN</b>	<b>288</b>
21.1.1	A Note on Responsible Disclosure	288
21.1.2	The Captcha Breaker Directory Structure	290
21.1.3	Automatically Downloading Example Images	291
21.1.4	Annotating and Creating Our Dataset	292
21.1.5	Preprocessing the Digits	297
21.1.6	Training the Captcha Breaker	299
21.1.7	Testing the Captcha Breaker	303
<b>21.2</b>	<b>Summary</b>	<b>305</b>
<b>22</b>	<b>Case Study: Smile Detection</b>	<b>307</b>
<b>22.1</b>	<b>The SMILES Dataset</b>	<b>307</b>

22.2	Training the Smile CNN	308
22.3	Running the Smile CNN in Real-time	313
22.4	Summary	316
23	Your Next Steps .....	319
23.1	So, What's Next?	319

## Companion Website

Thank you for picking up a copy of *Deep Learning for Computer Vision with Python!* To accompany this book I have created a companion website which includes:

- Up-to-date installation instructions on how to configure your development environment
- Instructions on how to use the pre-configured Ubuntu VirtualBox virtual machine and Amazon Machine Image (AMI)
- Supplementary material that I could not fit into this book
- Frequently Asked Questions (FAQs) and their suggested fixes and solutions

Additionally, you can use the "Issues" feature inside the companion website to report any bugs, typos, or problems you encounter while reading through the book. I don't expect many problems, however, this is a brand new book so myself and other readers would appreciate reporting any issues you run into. Your future, I can keep the book updated and buy me.

To create your companion website account, just use this link:

<http://bit.ly/2912933>

Take a second to create your account now so you'll have access to the supplementary materials as you work through the book.