

# Contents

<b>1</b>	<b>Introduction</b>	<b>13</b>
<b>2</b>	<b>Data Augmentation</b>	<b>15</b>
<b>2.1</b>	<b>What Is Data Augmentation?</b>	<b>15</b>
<b>2.2</b>	<b>Visualizing Data Augmentation</b>	<b>16</b>
<b>2.3</b>	<b>Comparing Training With and Without Data Augmentation</b>	<b>19</b>
2.3.1	The Flowers-17 Dataset	19
2.3.2	Aspect-aware Preprocessing	20
2.3.3	Flowers-17: No Data Augmentation	23
2.3.4	Flowers-17: With Data Augmentation	27
<b>2.4</b>	<b>Summary</b>	<b>31</b>
<b>3</b>	<b>Networks as Feature Extractors</b>	<b>33</b>
<b>3.1</b>	<b>Extracting Features with a Pre-trained CNN</b>	<b>34</b>
3.1.1	What Is HDF5?	35
3.1.2	Writing Features to an HDF5 Dataset	36
<b>3.2</b>	<b>The Feature Extraction Process</b>	<b>39</b>
3.2.1	Extracting Features From Animals	43
3.2.2	Extracting Features From CALTECH-101	44
3.2.3	Extracting Features From Flowers-17	44
<b>3.3</b>	<b>Training a Classifier on Extracted Features</b>	<b>45</b>
3.3.1	Results on Animals	47
3.3.2	Results on CALTECH-101	47
3.3.3	Results on Flowers-17	48
<b>3.4</b>	<b>Summary</b>	<b>48</b>

<b>4</b>	<b>Understanding rank-1 &amp; rank-5 Accuracies</b>	<b>51</b>
<b>4.1</b>	<b>Ranked Accuracy</b>	<b>51</b>
4.1.1	Measuring rank-1 and rank-5 Accuracies	53
4.1.2	Implementing Ranked Accuracy	54
4.1.3	Ranked Accuracy on Flowers-17	56
4.1.4	Ranked Accuracy on CALTECH-101	56
<b>4.2</b>	<b>Summary</b>	<b>56</b>
<b>5</b>	<b>Fine-tuning Networks</b>	<b>59</b>
<b>5.1</b>	<b>Transfer Learning and Fine-tuning</b>	<b>59</b>
5.1.1	Indexes and Layers	62
5.1.2	Network Surgery	63
5.1.3	Fine-tuning, from Start to Finish	65
<b>5.2</b>	<b>Summary</b>	<b>71</b>
<b>6</b>	<b>Improving Accuracy with Ensembles</b>	<b>73</b>
<b>6.1</b>	<b>Ensemble Methods</b>	<b>73</b>
6.1.1	Jensen's Inequality	74
6.1.2	Constructing an Ensemble of CNNs	75
6.1.3	Evaluating an Ensemble	79
<b>6.2</b>	<b>Summary</b>	<b>82</b>
<b>7</b>	<b>Advanced Optimization Methods</b>	<b>85</b>
<b>7.1</b>	<b>Adaptive Learning Rate Methods</b>	<b>85</b>
7.1.1	Adagrad	86
7.1.2	Adadelta	86
7.1.3	RMSprop	87
7.1.4	Adam	87
7.1.5	Nadam	88
<b>7.2</b>	<b>Choosing an Optimization Method</b>	<b>88</b>
7.2.1	Three Methods You Should Learn how to Drive: SGD, Adam, and RMSprop	88
<b>7.3</b>	<b>Summary</b>	<b>89</b>
<b>8</b>	<b>Optimal Pathway to Apply Deep Learning</b>	<b>91</b>
<b>8.1</b>	<b>A Recipe for Training</b>	<b>91</b>
<b>8.2</b>	<b>Transfer Learning or Train from Scratch</b>	<b>95</b>
<b>8.3</b>	<b>Summary</b>	<b>96</b>
<b>9</b>	<b>Working with HDF5 and Large Datasets</b>	<b>97</b>
<b>9.1</b>	<b>Downloading Kaggle: Dogs vs. Cats</b>	<b>97</b>
<b>9.2</b>	<b>Creating a Configuration File</b>	<b>98</b>
9.2.1	Your First Configuration File	99
<b>9.3</b>	<b>Building the Dataset</b>	<b>100</b>
<b>9.4</b>	<b>Summary</b>	<b>104</b>

<b>10</b>	<b>Competing in Kaggle: Dogs vs. Cats</b>	<b>105</b>
<b>10.1</b>	<b>Additional Image Preprocessors</b>	<b>105</b>
10.1.1	Mean Preprocessing	106
10.1.2	Patch Preprocessing	107
10.1.3	Crop Preprocessing	109
<b>10.2</b>	<b>HDF5 Dataset Generators</b>	<b>111</b>
<b>10.3</b>	<b>Implementing AlexNet</b>	<b>114</b>
<b>10.4</b>	<b>Training AlexNet on Kaggle: Dogs vs. Cats</b>	<b>119</b>
<b>10.5</b>	<b>Evaluating AlexNet</b>	<b>122</b>
<b>10.6</b>	<b>Obtaining a Top-5 Spot on the Kaggle Leaderboard</b>	<b>125</b>
10.6.1	Extracting Features Using ResNet	126
10.6.2	Training a Logistic Regression Classifier	129
<b>10.7</b>	<b>Summary</b>	<b>131</b>
<b>11</b>	<b>GoogLeNet</b>	<b>133</b>
<b>11.1</b>	<b>The Inception Module (and its Variants)</b>	<b>134</b>
11.1.1	Inception	134
11.1.2	Miniception	135
<b>11.2</b>	<b>MiniGoogLeNet on CIFAR-10</b>	<b>136</b>
11.2.1	Implementing MiniGoogLeNet	137
11.2.2	Training and Evaluating MiniGoogLeNet on CIFAR-10	142
11.2.3	MiniGoogLeNet: Experiment #1	145
11.2.4	MiniGoogLeNet: Experiment #2	146
11.2.5	MiniGoogLeNet: Experiment #3	147
<b>11.3</b>	<b>The Tiny ImageNet Challenge</b>	<b>148</b>
11.3.1	Downloading Tiny ImageNet	149
11.3.2	The Tiny ImageNet Directory Structure	149
11.3.3	Building the Tiny ImageNet Dataset	150
<b>11.4</b>	<b>DeeperGoogLeNet on Tiny ImageNet</b>	<b>155</b>
11.4.1	Implementing DeeperGoogLeNet	155
11.4.2	Training DeeperGoogLeNet on Tiny ImageNet	163
11.4.3	Creating the Training Script	163
11.4.4	Creating the Evaluation Script	165
11.4.5	DeeperGoogLeNet Experiments	167
<b>11.5</b>	<b>Summary</b>	<b>170</b>
<b>12</b>	<b>ResNet</b>	<b>173</b>
<b>12.1</b>	<b>ResNet and the Residual Module</b>	<b>173</b>
12.1.1	Going Deeper: Residual Modules and Bottlenecks	174
12.1.2	Rethinking the Residual Module	176
<b>12.2</b>	<b>Implementing ResNet</b>	<b>177</b>
<b>12.3</b>	<b>ResNet on CIFAR-10</b>	<b>182</b>
12.3.1	Training ResNet on CIFAR-10 With the ctrl + c Method	183
12.3.2	ResNet on CIFAR-10: Experiment #2	187

<b>12.4</b>	<b>Training ResNet on CIFAR-10 with Learning Rate Decay</b>	<b>190</b>
<b>12.5</b>	<b>ResNet on Tiny ImageNet</b>	<b>194</b>
12.5.1	Updating the ResNet Architecture	195
12.5.2	Training ResNet on Tiny ImageNet With the ctrl + c Method	196
12.5.3	Training ResNet on Tiny ImageNet with Learning Rate Decay	200
<b>12.6</b>	<b>Summary</b>	<b>204</b>
<b>13</b>	<b>Fundamentals of Object Detection</b>	<b>207</b>
<b>13.1</b>	<b>A Brief Object Detection History Lesson</b>	<b>207</b>
13.1.1	Haar Cascades and HOG + Linear SVM	208
<b>13.2</b>	<b>The Object Detection Pipeline</b>	<b>208</b>
13.2.1	Sliding Windows	209
13.2.2	Image Pyramids	211
13.2.3	Batch Processing	212
13.2.4	Non-maxima Suppression	213
<b>13.3</b>	<b>Implementing Object Detection With a CNN</b>	<b>214</b>
<b>13.4</b>	<b>Simple Object Detection Results</b>	<b>218</b>
<b>13.5</b>	<b>Summary</b>	<b>218</b>
<b>14</b>	<b>DeepDream</b>	<b>221</b>
<b>14.1</b>	<b>What Is DeepDream?</b>	<b>221</b>
<b>14.2</b>	<b>Implementing DeepDream</b>	<b>222</b>
<b>14.3</b>	<b>DeepDream Results</b>	<b>229</b>
<b>14.4</b>	<b>Summary</b>	<b>230</b>
<b>15</b>	<b>Neural Style Transfer</b>	<b>231</b>
<b>15.1</b>	<b>The Neural Style Transfer Algorithm</b>	<b>231</b>
15.1.1	Content Loss	232
15.1.2	Style Loss	233
15.1.3	Total-variation Loss	233
15.1.4	Combining the Loss Functions	233
<b>15.2</b>	<b>Implementing Neural Style Transfer</b>	<b>234</b>
15.2.1	Implementing the Style Transfer Driver Script	234
15.2.2	Implementing the Style Transfer Class	235
<b>15.3</b>	<b>Neural Style Transfer Results</b>	<b>242</b>
<b>15.4</b>	<b>Summary</b>	<b>243</b>
<b>16</b>	<b>Generative Adversarial Networks (GANs)</b>	<b>245</b>
<b>16.1</b>	<b>What Are GANs?</b>	<b>245</b>
16.1.1	The General Training Procedure	246
16.1.2	Guidelines and Best Practices When Training	246
<b>16.2</b>	<b>Implementing a DCGAN</b>	<b>247</b>
<b>16.3</b>	<b>Training a DCGAN</b>	<b>251</b>
<b>16.4</b>	<b>GAN Results</b>	<b>256</b>

<b>16.5</b>	<b>Summary</b>	<b>257</b>
<b>17</b>	<b>Image Super Resolution</b>	<b>259</b>
<b>17.1</b>	<b>Understanding SRCNNs</b>	<b>259</b>
<b>17.2</b>	<b>Implementing SRCNNs</b>	<b>260</b>
17.2.1	Starting the Project	260
17.2.2	Building the Dataset	262
17.2.3	The SRCNN Architecture	265
17.2.4	Training the SRCNN	266
17.2.5	Increasing Image Resolution With SRCNNs	268
<b>17.3</b>	<b>Super Resolution Results</b>	<b>270</b>
<b>17.4</b>	<b>Summary</b>	<b>271</b>
<b>18</b>	<b>Where to Now?</b>	<b>273</b>
<b>18.1</b>	<b>What's Next?</b>	<b>274</b>

Thank you for picking up a copy of *Deep Learning for Computer Vision with Python*. As part of 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 inside 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 when working 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. From there, I can keep the book updated and bug free.

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

<http://mng.bz.com/2019>

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