

A Hands-On Guide to Machine Learning, Deep Learning,
AI, and Natural Language Processing

MACHINE LEARNING IS FUN!



ADAM GEITGEY

SECOND EDITION

Machine Learning Is Fun!

A Hands-On Guide to Machine Learning, AI, and Natural Language
Processing

Adam Geitgey

Second Edition, Version 2.0.0

Contents

1	Welcome	11
2	What Is Machine Learning?	12
	Types of Machine Learning Algorithms	14
	Supervised Learning	14
	Unsupervised Learning	16
	More Esoteric Types of Machine Learning Algorithms	18
	Machine Learning Is Not the Same as Human Learning	18
	Let's Write That Program!	19
	The Power of Machine Learning	24
	Training: Finding Weights Automatically	25
	How Gradient Descent Works	26
	The Limitations of a Linear Model	29
3	The Machine Learning Toolbox	30
	Choosing the Right Operating System	30
	The Python 3 Programming Language	32
	The Python Library Ecosystem	32
	NumPy	33
	Pandas	33
	Matplotlib	34

Contents

Scikit-learn	34
TensorFlow	35
Keras	35
PyTorch	36
Dlib	36
OpenCV	37
spaCy	37
fastText	38
Rasa NLU	38
LIME	38
Speeding Up Machine Learning with a GPU	38
Running Models in the Cloud with Google Colab	40
4 What Is a Neural Network?	41
Stacking More Layers	51
Neural Networks Require Feature Scaling	53
How to Train a Neural Network	54
5 What Is Deep Learning?	58
Does Deep Learning Count as AI?	59
Why Deep Learning Is so Useful	60
6 Convolutional Neural Networks	62
Simple Image Recognition	63
Recognizing More Complex Images	68
Convolution for Translational Invariance	72
How Convolution Works	74
How to Design a Neural Network for Maximum Accuracy	81
VGG—The Foundational Design	82

Contents

ResNet—A Solution for Deeper Networks	82
Inception-Based Designs	83
ResNeXt: The Best of ResNet and Inception-Based Designs	83
Mobile-Focused Designs	84
7 Transfer Learning	85
How to Transfer Knowledge Between Networks	86
Making Predictions with Transfer Learning	89
Where to Get a Pre-Trained Neural Network	89
When to Use Transfer Learning	90
Other Forms of Transfer Learning	91
Retraining Convolutional Layers	91
Fine-Tuning the Entire Neural Network	92
8 Image Segmentation with Mask R-CNN	94
The Uses of Image Segmentation	95
Building an Image Segmentation Model	98
Adding Object Localization to Our Model	98
Classifying the Detected Objects	107
Adding Object Mask Generation to Our Model	108
The Complete Model	110
Implementing Mask R-CNN	112
The Magic of Image Segmentation	112
9 Face Recognition with Deep Learning	118
Creating a Face Recognition Pipeline	119
Face Recognition—Step by Step	121
Step 1: Finding All the Faces	121
Step 2: Posing and Projecting Faces	126

Contents

Step 3: Encoding Faces	128
Step 4: Finding the Person’s Identity from the Encoding . . .	132
Running the Face Recognition Pipeline	133
The Hidden Risks of Face Recognition Models	134
10 Generative Adversarial Networks	136
The Goal of Generative Models	138
How DCGANs Work	140
Exploring the Space Between	144
Getting Data	146
Setting Up the DCGAN	147
GANs Grow Up	152
11 Speech Recognition with Deep Learning	158
Speech Recognition Is Still Difficult	159
Turning Sounds into Bits	160
Sampling Is More Accurate Than It Sounds	163
Pre-Processing Sampled Sound Data	164
Recognizing Characters From Short Sounds	166
Improving Accuracy by Adding a Language Model	169
Building a Custom Speech Recognition System Is Still Difficult . .	170
12 Natural Language Processing	173
Extracting Meaning from Text Is Hard	174
Building an NLP Pipeline, Step by Step	174
Step 1: Sentence Segmentation	175
Step 2: Word Tokenization	176
Step 3: Predicting Parts of Speech for Each Token	176
Step 4: Text Lemmatization	177

Contents

Step 5: Identifying Stop Words	178
Step 6: Dependency Parsing	179
Step 6b: Finding Noun Phrases	181
Step 7: Named Entity Recognition (NER)	182
Step 8: Coreference Resolution	184
Putting It All Together	185
13 Understanding Language with Text Classification	187
Working Bottom Up Instead of Top Down	187
Using Classification Models to Extract Meaning	188
Why Does This Work? It Seems Too Simple!	190
Uses of Text Classification	192
14 Explaining Model Predictions with LIME	196
The Problem with Black Box Models	197
Looking Inside a Complex Machine Learning Model	198
We Need a Stunt Double!	201
LIME Step by Step	204
Step 1: Make the Initial Prediction with fastText	205
Step 2: Create Variations of the Review	205
Step 4: Train the Stand-In Model	206
Step 5: Examine the Stand-In Model to Explain the Original Prediction	207
15 How Digital Assistants and Chatbots Understand People	211
The Chat Bot NLP Pipeline	212
Step 1: Intent Detection Model	214
Step 2: Entity Recognition	215
Step 3: Clarifying Missing Information	217

Contents

Step 4: Query Auto-Correction	219
Step 5: Execute the Requested Action	220
The Bot’s Main Loop	221
Plugging a Bot into Slack or Facebook	221
Making a Chat Bot Useful	221
16 Modeling Text with Recurrent Neural Networks	223
Learning from Sequences	225
Using a Recurrent Model in Practice	227
Uses of Recurrent Neural Networks	228
Generating an Endless Story	230
17 Sequence Models and Automated Language Translation	235
Brute Force Machine Translation	236
Statistical Machine Translation	238
Step 1: Break the original sentence into chunks	239
Step 2: Find all possible translations for each chunk	239
Step 3: Generate all possible sentences and find the most likely one	240
Statistical Machine Translation Was a Huge Milestone	242
The Limitations of Statistical Machine Translation	243
Making Computers Translate Better—Without All Those Expensive People	243
Let’s Translate!	246
Dealing with Long Sentences	249
Building your own Translation System	250
The Power of Sequence-to-Sequence Models	251

Contents

18 Cutting Edge Text Generation with GPT-2 and Grover	255
The Entire Universe, One Word at a Time	256
The GPT-2 Model	257
Grover—Extending GPT-2 to Control the Output	259
How Original Is the Generated Text?	262
New Tools Have New Uses We Can't Yet Imagine	262
19 Using Machine Learning for Evil	264
Machine Learning and Hidden Bias	265
20 Tricking Neural Networks	269
Neural Nets as Security Guards	270
Neural Nets Are Not as Reliable as They Seem	273
How to Trick a Neural Network	277
What Can We Do with a Hacked Image?	281
How Can We Protect Ourselves Against These Attacks?	283
21 How to Apply Machine Learning to Your Business	285
When Is Machine Learning the Right Choice?	285
Machine Learning Is Not Magic	286
Machine Learning Is Best at Automating Repetitive Tasks	287
Machine Learning Is Data Hungry	288
The Model / Complexity Continuum	289
Don't Try to Beat Human Performance	290
How to Build Your Business in a Way that Makes Machine Learning Possible	291
How to Get Your Company Into Shape	293
22 The End	297