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

MACHINE LEARNING IS FUN!

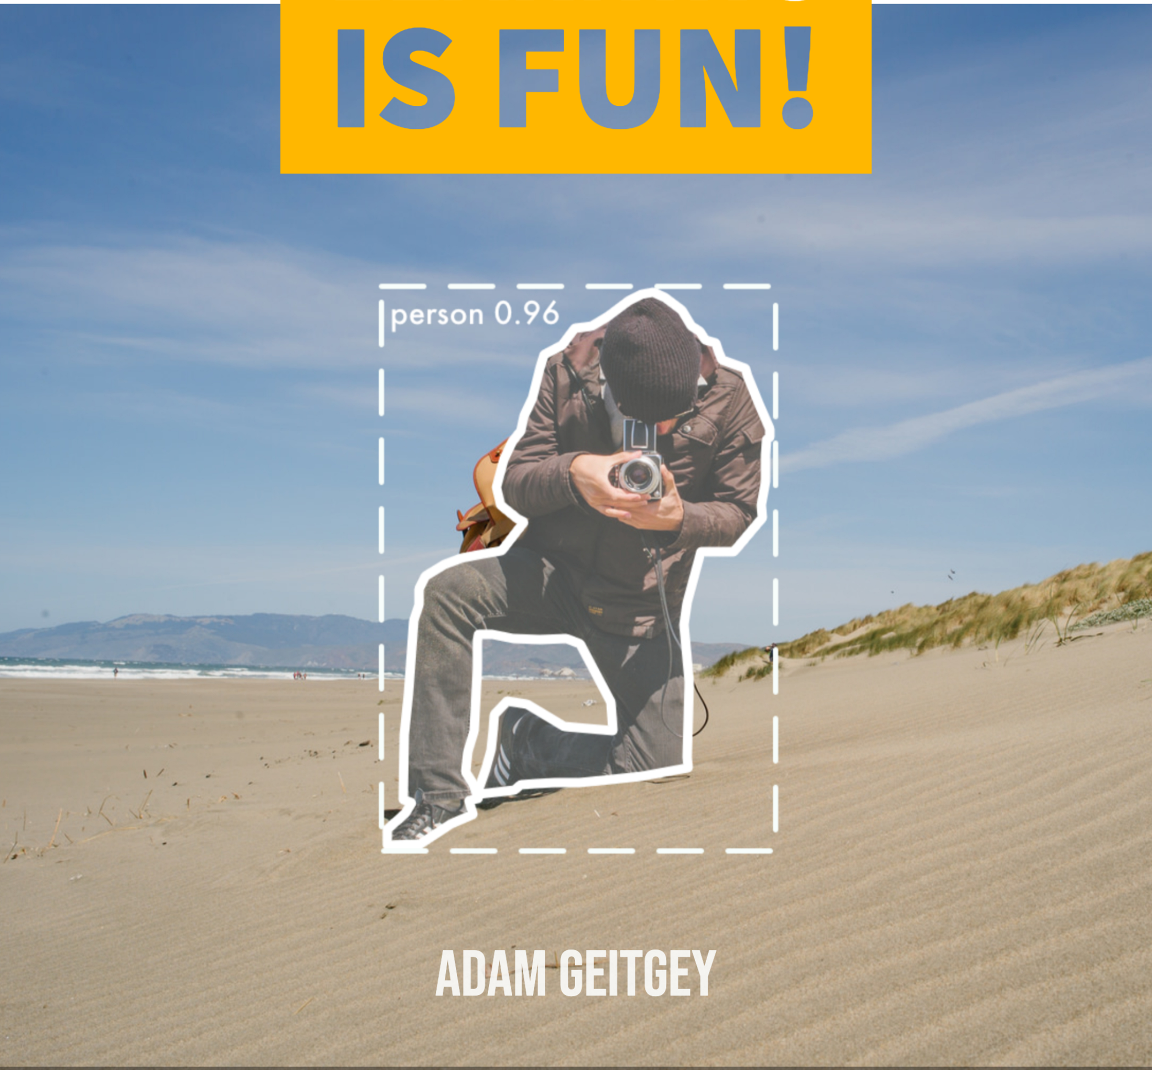

ADAM GEITGEY

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 Second Edition **16**

TensorFlow 2.0 and Keras Updates 16

New Chapters 16

2 Welcome **18**

3 What Is Machine Learning? **19**

Types of Machine Learning Algorithms 21

Supervised Learning 21

Unsupervised Learning 23

More Esoteric Types of Machine Learning Algorithms 25

Machine Learning Is Not the Same as Human Learning 25

Let's Write That Program! 26

The Power of Machine Learning 31

Training: Finding Weights Automatically 32

How Gradient Descent Works 33

The Limitations of a Linear Model 36

4 The Machine Learning Toolbox **37**

Choosing the Right Operating System 37

The Python 3 Programming Language 39

The Python Library Ecosystem 39
- NumPy 40
- Pandas 40
- Matplotlib 41
- Scikit-learn 41
- TensorFlow 42
- Keras 42
- PyTorch 43
- Dlib 43
- OpenCV 44
- spaCy 44
- fastText 45
- Rasa NLU 45
- LIME 45

Speeding Up Machine Learning with a GPU 45

Running Models in the Cloud with Google Colab 47

5 Project: Training a Regression Model Using Linear Regression 48

Examining the Training Data 48

Creating the Model 49

Using the Model 53

Extra Credit—Trying Another Algorithm 55

6 What Is a Neural Network? 57

Stacking More Layers 67

Neural Networks Require Feature Scaling 69

How to Train a Neural Network 70

7 Project: Training a Regression Model Using a Neural Network 74
Training the Model 74
Using the Model 84

8 What Is Deep Learning? 88
Does Deep Learning Count as AI? 89
Why Deep Learning Is so Useful 90

9 Convolutional Neural Networks 92
Simple Image Recognition 93
Recognizing More Complex Images 98
Convolution for Translational Invariance 102
How Convolution Works 104
How to Design a Neural Network for Maximum Accuracy 111
VGG—The Foundational Design 112
ResNet—A Solution for Deeper Networks 112
Inception-Based Designs 113
ResNeXt: The Best of ResNet and Inception-Based Designs 113
Mobile-Focused Designs 114

10 Project: Building a Bird Classifier 115
Getting Training Data 115
Coding the Model 117
Defining the Neural Network 120
Testing the Neural Network 126
How accurate is 95 percent accurate? 129

11 Transfer Learning 137
How to Transfer Knowledge Between Networks 138
Making Predictions with Transfer Learning 141

Where to Get a Pre-Trained Neural Network 141
When to Use Transfer Learning 142
Other Forms of Transfer Learning 143
Retraining Convolutional Layers 143
Fine-Tuning the Entire Neural Network 144

12 Project: Building a Bird Classifier with Transfer Learning 146
Extracting Features with a Pre-Trained CNN 147
Training Our New Classifier 152
Using the Model to Make Predictions 155
Checking Precision and Recall 158
Using Transfer Learning in Your Projects 163

13 Image Segmentation with Mask R-CNN 164
The Uses of Image Segmentation 165
Building an Image Segmentation Model 168
Adding Object Localization to Our Model 168
Classifying the Detected Objects 177
Adding Object Mask Generation to Our Model 178
The Complete Model 180
Implementing Mask R-CNN 182
The Magic of Image Segmentation 182

14 Project: Using a Pre-trained Mask R-CNN Model 188
Activating the Mask-RCNN Virtual Environment 189
Setting Up the Model 191
Counting People in a Line 192

15 Project: Training a Mask R-CNN Segmentation Model 205
Gathering and Annotating Training Data 207

Activating the Mask-RCNN Virtual Environment 210
Coding the Training Script 210
Parsing the Annotation Files 212
Running the Training Process 218
Using the Trained Model for Image Segmentation 222

16 Project: Detecting Available Parking Spaces with Mask R-CNN 228
Breaking Down the Problem 229
Activating the Mask-RCNN Virtual Environment 231
Writing the Code 231
Step 1: Detecting Parking Spaces in an Image 234
Step 2: Detecting Cars in an Image 239
Step 3: Detecting Empty Parking Spaces 244
Triggering a Detection 248
Running the Code 250

17 Face Recognition with Deep Learning 252
Creating a Face Recognition Pipeline 253
Face Recognition—Step by Step 255
Step 1: Finding All the Faces 255
Step 2: Posing and Projecting Faces 260
Step 3: Encoding Faces 262
Step 4: Finding the Person’s Identity from the Encoding . . 266
Running the Face Recognition Pipeline 267
The Hidden Risks of Face Recognition Models 268

18 Project: Recognizing Faces 270
Gathering Training Images 270
Detecting and Encoding Faces 271

19 Project: Doorbell Camera with Face Recognition 277
Detecting and Tracking Unknown People 278
Coding the Doorbell Camera 279
Running the Code 288

20 Project: Deploying the Doorbell Camera on Real Hardware 290
Choosing a Hardware Platform 291
What to Buy 294
Downloading the Jetson Nano Software 296
Plugging Everything In 296
First Boot and User Account Configuration 300
Installing Required Python Libraries 301
Adapting the Doorbell Camera Code 303
Running the Face Recognition Doorbell Camera Code 305

21 Generative Adversarial Networks 307
The Goal of Generative Models 309
How DCGANs Work 311
Exploring the Space Between 315
Getting Data 317
Setting Up the DCGAN 318
GANs Grow Up 323

22 Speech Recognition with Deep Learning 329
Speech Recognition Is Still Difficult 330
Turning Sounds into Bits 331
Sampling Is More Accurate Than It Sounds 334
Pre-Processing Sampled Sound Data 335
Recognizing Characters From Short Sounds 337

Improving Accuracy by Adding a Language Model 340
Building a Custom Speech Recognition System Is Still Difficult . 341

23 Project: Recognizing Speech with DeepSpeech 344
Preparing Audio to Be Processed with Deep Speech 345
Downloading the Pre-Trained Model Data 347
Recognizing Audio with the Pre-Trained Model 348
Running the Code 353
When to Retrain with Custom Training Data 353

24 Natural Language Processing 356
Extracting Meaning from Text Is Hard 357
Building an NLP Pipeline, Step by Step 357
Step 1: Sentence Segmentation 358
Step 2: Word Tokenization 359
Step 3: Predicting Parts of Speech for Each Token 359
Step 4: Text Lemmatization 360
Step 5: Identifying Stop Words 361
Step 6: Dependency Parsing 362
Step 6b: Finding Noun Phrases 364
Step 7: Named Entity Recognition (NER) 365
Step 8: Coreference Resolution 367
Putting It All Together 368

25 Project: Extracting Facts with an NLP Pipeline 370
Extracting Entities from Text 370
Building a Data Scrubber 373
Extracting Facts from Text 376
Extracting Noun Chucks 379

26 Project: Using NLP to Extract Podcast Appearances 383

Building Our Data Set 384

Parsing Podcast Feeds to Create the Dataset 385

Extracting Data from Plain Text with NLP 387

Writing the Program 388

27 Understanding Language with Text Classification 394

Working Bottom Up Instead of Top Down 394

Using Classification Models to Extract Meaning 395

Why Does This Work? It Seems Too Simple! 397

Uses of Text Classification 399

28 Project: Training and Using a Text Classifier 403

Getting the Training Data 403

Formatting the Training Data 405

Normalizing the Data 405

Splitting the Data into a Training Set and a Test Set 406

Writing the Data Preparation Script 406

Training the Model 409

Testing the Model 410

Iterating on the Model 412

Using the Model in a Program 414

29 Explaining Model Predictions with LIME 418

The Problem with Black Box Models 419

Looking Inside a Complex Machine Learning Model 420

We Need a Stunt Double! 423

LIME Step by Step 426

Step 1: Make the Initial Prediction with fastText 427

Step 2: Create Variations of the Review 427
Step 4: Train the Stand-In Model 428
Step 5: Examine the Stand-In Model to Explain the Original Prediction 429

30 Project: Using LIME to Explain Text Classifications 433

31 How Digital Assistants and Chatbots Understand People 441
The Chat Bot NLP Pipeline 442
Step 1: Intent Detection Model 444
Step 2: Entity Recognition 445
Step 3: Clarifying Missing Information 447
Step 4: Query Auto-Correction 449
Step 5: Execute the Requested Action 450
The Bot's Main Loop 451
Plugging a Bot into Slack or Facebook 451
Making a Chat Bot Useful 451

32 Project: Building a Chatbot 453
Creating the Training Data 453
Defining the Model Configuration 456
Training the Model 457
Testing the Model 459
Coding the Bot 460
Running the Bot 465

33 Modeling Text with Recurrent Neural Networks 467
Learning from Sequences 469
Using a Recurrent Model in Practice 471
Uses of Recurrent Neural Networks 472

Generating an Endless Story 474

34 Project: Generating Text with a Recurrent Model 479

Building the Model 479

Training the Model 491

35 Sequence Models and Automated Language Translation 495

Brute Force Machine Translation 496

Statistical Machine Translation 498

Step 1: Break the original sentence into chunks 499

Step 2: Find all possible translations for each chunk 499

Step 3: Generate all possible sentences and find the most likely one 500

Statistical Machine Translation Was a Huge Milestone . . . 502

The Limitations of Statistical Machine Translation 503

Making Computers Translate Better—Without All Those Expensive People 503

Let's Translate! 506

Dealing with Long Sentences 509

Building your own Translation System 510

The Power of Sequence-to-Sequence Models 511

36 Cutting Edge Text Generation with GPT-2 and Grover 515

The Entire Universe, One Word at a Time 516

The GPT-2 Model 517

Grover—Extending GPT-2 to Control the Output 519

How Original Is the Generated Text? 522

New Tools Have New Uses We Can't Yet Imagine 522

37 Project: Generating Realistic News Articles 524
Writing the Code 526
Configuration 526
Download Grover Code and Install Requirements 527
Download Grover Pre-Trained “Mega” Model 528
Generate News Stores 530
Running It 540
Evaluating the Results 541

38 Using Machine Learning for Evil 544
Machine Learning and Hidden Bias 545

39 Project: How to break CAPTCHAs with CNNs 549
Picking a Target 550
Examining the CAPTCHA 551
Creating the Dataset 552
Breaking Up Words 553
Writing the Code 558
Building and Training the Neural Network 565
Using the Trained Model to Solve CAPTCHAs 572

40 Tricking Neural Networks 574
Neural Nets as Security Guards 575
Neural Nets Are Not as Reliable as They Seem 578
How to Trick a Neural Network 582
What Can We Do with a Hacked Image? 586
How Can We Protect Ourselves Against These Attacks? 588

41 Project: Creating Adversarial Images 590
Validating the Starting Image 590

Creating the Adversarial Image 593

42 How to Apply Machine Learning to Your Business 602

When Is Machine Learning the Right Choice? 602

Machine Learning Is Not Magic 603

Machine Learning Is Best at Automating Repetitive Tasks 604

Machine Learning Is Data Hungry 605

The Model / Complexity Continuum 606

Don't Try to Beat Human Performance 607

How to Build Your Business in a Way that Makes Machine Learning Possible 608

How to Get Your Company Into Shape 610

43 The End 614