

Table of Contents

Preface	1
Chapter 1: Setting Up OpenCV	7
Technical requirements	8
What's new in OpenCV 4	9
Choosing and using the right setup tools	9
Installation on Windows	10
Using a ready-made OpenCV package	11
Building OpenCV from source	11
Installation on macOS	15
Using Homebrew with ready-made packages	15
Using Homebrew with your own custom packages	17
Installation on Debian, Ubuntu, Linux Mint, and similar systems	17
Using a ready-made OpenCV package	18
Building OpenCV from source	19
Installation on other Unix-like systems	21
Running samples	22
Finding documentation, help, and updates	23
Summary	24
Chapter 2: Handling Files, Cameras, and GUIs	25
Technical requirements	26
Basic I/O scripts	26
Reading/writing an image file	26
Converting between an image and raw bytes	29
Accessing image data with numpy.array	32
Reading/writing a video file	34
Capturing camera frames	36
Displaying an image in a window	38
Displaying camera frames in a window	39
Project Cameo (face tracking and image manipulation)	41
Cameo – an object-oriented design	42
Abstracting a video stream with managers.CaptureManager	42
Abstracting a window and keyboard with managers.WindowManager	48
Applying everything with cameo.Cameo	49
Summary	52
Chapter 3: Processing Images with OpenCV	53
Technical requirements	53
Converting images between different color models	54
Light is not paint	55

Exploring the Fourier transform	55
HPFs and LPFs	56
Creating modules	60
Edge detection	60
Custom kernels – getting convoluted	62
Modifying the application	64
Edge detection with Canny	66
Contour detection	68
Bounding box, minimum area rectangle, and minimum enclosing circle	69
Convex contours and the Douglas-Peucker algorithm	73
Detecting lines, circles, and other shapes	75
Detecting lines	76
Detecting circles	77
Detecting other shapes	78
Summary	79
Chapter 4: Depth Estimation and Segmentation	81
Technical requirements	82
Creating modules	82
Capturing frames from a depth camera	83
Converting 10-bit images to 8-bit	85
Creating a mask from a disparity map	88
Modifying the application	89
Depth estimation with a normal camera	92
Foreground detection with the GrabCut algorithm	99
Image segmentation with the Watershed algorithm	103
Summary	106
Chapter 5: Detecting and Recognizing Faces	107
Technical requirements	108
Conceptualizing Haar cascades	108
Getting Haar cascade data	109
Using OpenCV to perform face detection	110
Performing face detection on a still image	111
Performing face detection on a video	113
Performing face recognition	117
Generating the data for face recognition	117
Recognizing faces	119
Loading the training data for face recognition	120
Performing face recognition with Eigenfaces	122
Performing face recognition with Fisherfaces	124
Performing face recognition with LBPH	125
Discarding results based on the confidence score	125
Swapping faces in the infrared	126
Modifying the application's loop	127

Masking a copy operation	130
Summary	133
Chapter 6: Retrieving Images and Searching Using Image Descriptors	135
Technical requirements	136
Understanding types of feature detection and matching	136
Defining features	137
Detecting Harris corners	137
Detecting DoG features and extracting SIFT descriptors	140
Anatomy of a keypoint	143
Detecting Fast Hessian features and extracting SURF descriptors	143
Using ORB with FAST features and BRIEF descriptors	145
FAST	146
BRIEF	147
Brute-force matching	147
Matching a logo in two images	148
Filtering matches using K-Nearest Neighbors and the ratio test	151
Matching with FLANN	155
Performing homography with FLANN-based matches	160
A sample application – tattoo forensics	164
Saving image descriptors to file	164
Scanning for matches	166
Summary	169
Chapter 7: Building Custom Object Detectors	171
Technical requirements	172
Understanding HOG descriptors	172
Visualizing HOG	173
Using HOG to describe regions of an image	175
Understanding NMS	176
Understanding SVMs	177
Detecting people with HOG descriptors	178
Creating and training an object detector	181
Understanding BoW	182
Applying BoW to computer vision	183
k-means clustering	184
Detecting cars	184
Combining an SVM with a sliding window	193
Detecting a car in a scene	194
Saving and loading a trained SVM	202
Summary	203
Chapter 8: Tracking Objects	205
Technical requirements	206
Detecting moving objects with background subtraction	206

Implementing a basic background subtractor	208
Using a MOG background subtractor	211
Using a KNN background subtractor	216
Using GMG and other background subtractors	218
Tracking colorful objects using MeanShift and CamShift	221
Planning our MeanShift sample	222
Calculating and back-projecting color histograms	223
Understanding the parameters of cv2.calcHist	226
Understanding the parameters of cv2.calcBackProject	227
Implementing the MeanShift example	228
Using CamShift	231
Finding trends in motion using the Kalman filter	232
Understanding the predict and update phases	233
Tracking a mouse cursor	234
Tracking pedestrians	237
Planning the flow of the application	237
Comparing the object-oriented and functional paradigms	238
Implementing the Pedestrian class	240
Implementing the main function	242
Considering the next steps	246
Summary	247
Chapter 9: Camera Models and Augmented Reality	249
Technical requirements	250
Understanding 3D image tracking and augmented reality	250
Understanding camera and lens parameters	253
Understanding cv2.solvePnP/Ransac	258
Implementing the demo application	261
Importing modules	261
Performing grayscale conversion	262
Performing 2D-to-3D spatial conversions	263
Implementing the application class	265
Initializing the tracker	266
Implementing the main loop	274
Tracking the image in 3D	275
Initializing and applying the Kalman filter	279
Drawing the tracking results and masking the tracked object	283
Running and testing the application	286
Improving the 3D tracking algorithm	291
Summary	292
Chapter 10: Introduction to Neural Networks with OpenCV	293
Technical requirements	294
Understanding ANNs	295
Understanding neurons and perceptrons	297
Understanding the layers of a neural network	298

Choosing the size of the input layer	299
Choosing the size of the output layer	299
Choosing the size of the hidden layer	299
Training a basic ANN in OpenCV	300
Training an ANN classifier in multiple epochs	303
Recognizing handwritten digits with an ANN	308
Understanding the MNIST database of handwritten digits	308
Choosing training parameters for the MNIST database	310
Implementing a module to train the ANN	310
Implementing a minimal test module	315
Implementing the main module	315
Trying to improve the ANN's training	321
Finding other potential applications	323
Using DNNs from other frameworks in OpenCV	324
Detecting and classifying objects with third-party DNNs	325
Detecting and classifying faces with third-party DNNs	329
Summary	336
Appendix A: Bending Color Space with the Curves Filter	339
Other Book You May Enjoy	349
Index	351
