2. Object Recognition using Neural Networks and Supervised Learning

This is the chapter where we start to combine robotics and artificial intelligence to accomplish some of the tasks we have laid out so carefully in the previous chapters. The subject of this chapter is image recognition – we will be teaching the robot to recognize what is a toy, and what is not a toy, so that it can then decide what to pick up and what to leave alone. We will be using convolutional neural networks as a machine learning tool to separate objects in images, recognize them, and locate them in the camera frame, so that the robot can then locate them.

In this chapter, we will cover the following topics:

- o The basics of image recognition: what is an image?
- o Artificial neural networks: what is a neuron, and how do they work? • Training neural networks using stochastic gradient descent
- o Image processing and dataset creation: training sets and test sets, multiplying data
- o Convolution: image processing for robots • The LeNet framework from Keras
- o Training and testing the network
- o Deploying our trained network on the Raspberry Pi 3 in the robot

Technical requirements

We will actually be able to accomplish all of these tasks without a robot, if yours cannot walk yet. We will, however, get better results if the camera is in the proper position on the robot. If you don't have a robot, you can still do all of these tasks with a laptop and a USB camera.

- Hardware: Laptop or portable computer, Raspberry Pi 3, USB camera
- Software:
  - Python
  - Open CV2
  - Keras (`https://keras.io/`)
  - Tensorflow (`https://www.tensorflow.org/install/`)
  - USB camera drivers installed on your laptop or Raspberry Pi

The source code for this chapter can be found at `https://github.com/PacktPublishing/ Artificial-Intelligence-for-Robotics/chapter4`.

Check out the following video to see the Code in Action:
`http://bit.ly/2wwM9Zl`