Programming Robots with ROS:

A Systematic Approach to Learning Robot Programming with ROS:

Artificial Intelligence for Robotic:

Chapter 1: Foundation for Advanced Robotics and AI 6

Chapter 2: Setting Up Your Robot 41

Deep Learning with Keras:

Further improving the simple net in Keras with dropout
Testing different optimizers in Keras
Increasing the number of epochs
Controlling the optimizer learning rate
Increasing the number of internal hidden neurons
Increasing the size of batch computation
Summarizing the experiments run for recognizing handwritten charts
Adopting regularization for avoiding overfitting
Hyperparameters tuning
Predicting output
A practical overview of backpropagation
Towards a deep learning approach
Summary

2. Keras Installation and API

Installing Keras
Step 1 — install some useful dependencies
Step 2 — install Theano
Step 3 — install TensorFlow
Step 4 — install Keras
Step 5 — testing Theano, TensorFlow, and Keras
Configuring Keras
Installing Keras on Docker
Installing Keras on Google Cloud ML
Installing Keras on Amazon AWS
Installing Keras on Microsoft Azure
Keras API
Getting started with Keras architecture
What is a tensor?
Composing models in Keras
Sequential composition
Functional composition
An overview of predefined neural network layers
Regular dense
Recurrent neural networks — simple, LSTM, and GRU
Convolutional and pooling layers
Regularization
Batch normalization
An overview of predefined activation functions
An overview of losses functions
An overview of metrics
An overview of optimizers
Some useful operations
Saving and loading the weights and the architecture of a model
Callbacks for customizing the training process
Checkpointing
Using TensorBoard and Keras
Using Quiver and Keras
Summary

3. Deep Learning with ConvNets
Deep convolutional neural network — DCNN
Local receptive fields
Shared weights and bias
Pooling layers
Max-pooling
Average pooling
ConvNets summary
An example of DCNN — LeNet
LeNet code in Keras
Understanding the power of deep learning
Recognizing CIFAR-10 images with deep learning
Improving the CIFAR-10 performance with deeper a network
Improving the CIFAR-10 performance with data augmentation
Predicting with CIFAR-10
Very deep convolutional networks for large-scale image recognition
Recognizing cats with a VGG-16 net
Utilizing Keras built-in VGG-16 net module
Recycling pre-built deep learning models for extracting features
Very deep inception-v3 net used for transfer learning
Summary
4. Generative Adversarial Networks and WaveNet
What is a GAN?
Some GAN applications
Deep convolutional generative adversarial networks
Keras adversarial GANs for forging MNIST
Keras adversarial GANs for forging CIFAR
WaveNet — a generative model for learning how to produce audio
Summary
5. Word Embeddings
Distributed representations
word2vec
The skip-gram word2vec model
The CBOW word2vec model
Extracting word2vec embeddings from the model
Using third-party implementations of word2vec
Exploring GloVe
Using pre-trained embeddings
Learn embeddings from scratch
Fine-tuning learned embeddings from word2vec
Fine-tune learned embeddings from GloVe
Look up embeddings
Summary
6. Recurrent Neural Network — RNN
SimpleRNN cells
SimpleRNN with Keras — generating text
RNN topologies
Vanishing and exploding gradients