


ADDISON WESLEY DATA & ANALYTICS SERIES



DEEP LEARNING ILLUSTRATED

A Visual, Interactive Guide to Artificial Intelligence



JON KROHN

with **GRANT BEYLEVELD** *and* **AGLAÉ BASSENS**

Praise for *Deep Learning Illustrated*

“Over the next few decades, artificial intelligence is poised to dramatically change almost every aspect of our lives, in large part due to today’s breakthroughs in deep learning. The authors’ clear visual style provides a comprehensive look at what’s currently possible with artificial neural networks as well as a glimpse of the magic that’s to come.”

—*Tim Urban, writer and illustrator of Wait But Why*

“This book is an approachable, practical, and broad introduction to deep learning, and the most beautifully illustrated machine learning book on the market.”

—*Dr. Michael Osborne, Dyson Associate Professor in Machine Learning, University of Oxford*

“This book should be the first stop for deep learning beginners, as it contains lots of concrete, easy-to-follow examples with corresponding tutorial videos and code notebooks. Strongly recommended.”

—*Dr. Chong Li, cofounder, Nakamoto & Turing Labs; adjunct professor, Columbia University*

“It’s hard to imagine developing new products today without thinking about enriching them with capabilities using machine learning. Deep learning in particular has many practical applications, and this book’s intelligible clear and visual approach is helpful to anyone who would like to understand what deep learning is and how it could impact your business and life for years to come.”

—*Helen Altshuler, engineering leader, Google*

Deep Learning Illustrated: A Visual, Interactive Guide to Artificial Intelligence

Table of Contents

Cover

Title Page

Copyright Page

Contents

Figures

Tables

Examples

Foreword

Preface

Acknowledgments

About the Authors

Part I: Introducing Deep Learning

1 Biological and Machine Vision

Biological Vision

Machine Vision

The Neocognitron

LeNet-5

The Traditional Machine Learning Approach

ImageNet and the ILSVRC

AlexNet

TensorFlow Playground

Table of Contents

Quick, Draw!

Summary

2 Human and Machine Language

Deep Learning for Natural Language Processing

Deep Learning Networks Learn Representations Automatically

Natural Language Processing

A Brief History of Deep Learning for NLP

Computational Representations of Language

One-Hot Representations of Words

Word Vectors

Word-Vector Arithmetic

word2viz

Localist Versus Distributed Representations

Elements of Natural Human Language

Google Duplex

Summary

3 Machine Art

A Boozy All-Nighter

Arithmetic on Fake Human Faces

Style Transfer: Converting Photos into Monet (and Vice Versa)

Make Your Own Sketches Photorealistic

Creating Photorealistic Images from Text

Image Processing Using Deep Learning

Summary

4 Game-Playing Machines

Deep Learning, AI, and Other Beasts

Artificial Intelligence

Machine Learning

Representation Learning

Artificial Neural Networks

Deep Learning

Table of Contents

Machine Vision

Natural Language Processing

Three Categories of Machine Learning Problems

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Deep Reinforcement Learning

Video Games

Board Games

AlphaGo

AlphaGo Zero

AlphaZero

Manipulation of Objects

Popular Deep Reinforcement Learning Environments

OpenAI Gym

DeepMind Lab

Unity ML-Agents

Three Categories of AI

Artificial Narrow Intelligence

Artificial General Intelligence

Artificial Super Intelligence

Summary

Part II: Essential Theory Illustrated

5 The (Code) Cart Ahead of the (Theory) Horse

Prerequisites

Installation

A Shallow Network in Keras

The MNIST Handwritten Digits

A Schematic Diagram of the Network

Loading the Data

Reformatting the Data

Table of Contents

Designing a Neural Network Architecture

Training a Deep Learning Model

Summary

6 Artificial Neurons Detecting Hot Dogs

Biological Neuroanatomy 101

The Perceptron

The Hot Dog / Not Hot Dog Detector

The Most Important Equation in This Book

Modern Neurons and Activation Functions

The Sigmoid Neuron

The Tanh Neuron

ReLU: Rectified Linear Units

Choosing a Neuron

Summary

Key Concepts

7 Artificial Neural Networks

The Input Layer

Dense Layers

A Hot Dog-Detecting Dense Network

Forward Propagation Through the First Hidden Layer

Forward Propagation Through Subsequent Layers

The Softmax Layer of a Fast Food-Classifying Network

Revisiting Our Shallow Network

Summary

Key Concepts

8 Training Deep Networks

Cost Functions

Quadratic Cost

Saturated Neurons

Cross-Entropy Cost

Optimization: Learning to Minimize Cost

Table of Contents

Gradient Descent

Learning Rate

Batch Size and Stochastic Gradient Descent

Escaping the Local Minimum

Backpropagation

Tuning Hidden-Layer Count and Neuron Count

An Intermediate Net in Keras

Summary

Key Concepts

9 Improving Deep Networks

Weight Initialization

Xavier Glorot Distributions

Unstable Gradients

Vanishing Gradients

Exploding Gradients

Batch Normalization

Model Generalization (Avoiding Overfitting)

L1 and L2 Regularization

Dropout

Data Augmentation

Fancy Optimizers

Momentum

Nesterov Momentum

AdaGrad

AdaDelta and RMSProp

Adam

A Deep Neural Network in Keras

Regression

TensorBoard

Summary

Key Concepts

Table of Contents

Part III: Interactive Applications of Deep Learning

10 Machine Vision

Convolutional Neural Networks

The Two-Dimensional Structure of Visual Imagery

Computational Complexity

Convolutional Layers

Multiple Filters

A Convolutional Example

Convolutional Filter Hyperparameters

Pooling Layers

LeNet-5 in Keras

AlexNet and VGGNet in Keras

Residual Networks

Vanishing Gradients: The Bête Noire of Deep CNNs

Residual Connections

ResNet

Applications of Machine Vision

Object Detection

Image Segmentation

Transfer Learning

Capsule Networks

Summary

Key Concepts

11 Natural Language Process

Preprocessing Natural Language Data

Tokenization

Converting All Characters to Lowercase

Removing Stop Words and Punctuation

Stemming

Handling n-grams

Preprocessing the Full Corpus

Creating Word Embeddings with word2vec

Table of Contents

The Essential Theory Behind word2vec

Evaluating Word Vectors

Running word2vec

Plotting Word Vectors

The Area under the ROC Curve

The Confusion Matrix

Calculating the ROC AUC Metric

Natural Language Classification with Familiar Networks

Loading the IMDb Film Reviews

Examining the IMDb Data

Standardizing the Length of the Reviews

Dense Network

Convolutional Networks

Networks Designed for Sequential Data

Recurrent Neural Networks

Long Short-Term Memory Units

Bidirectional LSTMs

Stacked Recurrent Models

Seq2seq and Attention

Transfer Learning in NLP

Non-sequential Architectures: The Keras Functional API

Summary

Key Concepts

12 Generative Adversarial Networks

Essential GAN Theory

The Quick, Draw! Dataset

The Discriminator Network

The Generator Network

The Adversarial Network

GAN Training

Summary

Key Concepts

Table of Contents

13 Deep Reinforcement Learning

Essential Theory of Reinforcement Learning

The Cart-Pole Game

Markov Decision Processes

The Optimal Policy

Essential Theory of Deep Q-Learning Networks

Value Functions

Q-Value Functions

Estimating an Optimal Q-Value

Defining a DQN Agent

Initialization Parameters

Building the Agents Neural Network Model

Remembering Gameplay

Training via Memory Replay

Selecting an Action to Take

Saving and Loading Model Parameters

Interacting with an OpenAI Gym Environment

Hyperparameter Optimization with SLM Lab

Agents Beyond DQN

Policy Gradients and the REINFORCE Algorithm

The Actor-Critic Algorithm

Summary

Key Concepts

Part IV: You and AI

14 Moving Forward with Your Own Deep Learning Projects

Ideas for Deep Learning Projects

Machine Vision and GANs

Natural Language Processing

Deep Reinforcement Learning

Converting an Existing Machine Learning Project

Resources for Further Projects

Table of Contents

Socially Beneficial Projects

The Modeling Process, Including Hyperparameter Tuning

Automation of Hyperparameter Search

Deep Learning Libraries

Keras and TensorFlow

PyTorch

MXNet, CNTK, Caffe, and So On

Software 2.0

Approaching Artificial General Intelligence

Summary

Part V: Appendices

Appendix A: Formal Neural Network Notation

Appendix B: Backpropagation

Appendix C: PyTorch

PyTorch Features

Autograd System

Define-by-Run Framework

PyTorch Versus TensorFlow

PyTorch in Practice

PyTorch Installation

The Fundamental Units Within PyTorch

Building a Deep Neural Network in PyTorch

Index