

Table of Contents

Preface	1
<hr/>	
Section 1: Getting Started with Deep Learning	
<hr/>	
Chapter 1: Introduction to Deep Learning	9
What is deep learning?	9
Biological and artificial neurons	11
ANN and its layers	13
Input layer	14
Hidden layer	14
Output layer	15
Exploring activation functions	15
The sigmoid function	15
The tanh function	16
The Rectified Linear Unit function	17
The leaky ReLU function	18
The Exponential linear unit function	20
The Swish function	20
The softmax function	21
Forward propagation in ANN	22
How does ANN learn?	25
Debugging gradient descent with gradient checking	32
Putting it all together	37
Building a neural network from scratch	37
Summary	40
Questions	41
Further reading	41
<hr/>	
Chapter 2: Getting to Know TensorFlow	43
What is TensorFlow?	43
Understanding computational graphs and sessions	44
Sessions	46
Variables, constants, and placeholders	47
Variables	47
Constants	49
Placeholders and feed dictionaries	49
Introducing TensorBoard	51
Creating a name scope	53
Handwritten digit classification using TensorFlow	55
Importing the required libraries	55

Loading the dataset	55
Defining the number of neurons in each layer	57
Defining placeholders	57
Forward propagation	58
Computing loss and backpropagation	59
Computing accuracy	60
Creating summary	61
Training the model	61
Visualizing graphs in TensorBoard	63
Introducing eager execution	68
Math operations in TensorFlow	69
TensorFlow 2.0 and Keras	73
Bonjour Keras	73
Defining the model	73
Defining a sequential model	74
Defining a functional model	74
Compiling the model	75
Training the model	76
Evaluating the model	76
MNIST digit classification using TensorFlow 2.0	76
Should we use Keras or TensorFlow?	77
Summary	78
Questions	79
Further reading	79

Section 2: Fundamental Deep Learning Algorithms

Chapter 3: Gradient Descent and Its Variants	83
Demystifying gradient descent	84
Performing gradient descent in regression	89
Importing the libraries	90
Preparing the dataset	90
Defining the loss function	91
Computing the gradients of the loss function	92
Updating the model parameters	94
Gradient descent versus stochastic gradient descent	96
Momentum-based gradient descent	99
Gradient descent with momentum	99
Nesterov accelerated gradient	101
Adaptive methods of gradient descent	103
Setting a learning rate adaptively using Adagrad	103
Doing away with the learning rate using Adadelta	106
Overcoming the limitations of Adagrad using RMSProp	110
Adaptive moment estimation	111
Adamax – Adam based on infinity-norm	113
Adaptive moment estimation with AMSGrad	116

Nadam – adding NAG to ADAM	117
Summary	120
Questions	121
Further reading	122
Chapter 4: Generating Song Lyrics Using RNN	123
Introducing RNNs	124
The difference between feedforward networks and RNNs	126
Forward propagation in RNNs	127
Backpropagating through time	130
Gradients with respect to the hidden to output weight, V	132
Gradients with respect to hidden to hidden layer weights, W	136
Gradients with respect to input to the hidden layer weight, U	142
Vanishing and exploding gradients problem	143
Gradient clipping	146
Generating song lyrics using RNNs	147
Implementing in TensorFlow	150
Data preparation	150
Defining the network parameters	153
Defining placeholders	154
Defining forward propagation	154
Defining BPTT	155
Start generating songs	156
Different types of RNN architectures	161
One-to-one architecture	161
One-to-many architecture	162
Many-to-one architecture	163
Many-to-many architecture	164
Summary	164
Questions	165
Further reading	165
Chapter 5: Improvements to the RNN	167
LSTM to the rescue	168
Understanding the LSTM cell	169
Forget gate	170
Input gate	171
Output gate	172
Updating the cell state	173
Updating hidden state	176
Forward propagation in LSTM	177
Backpropagation in LSTM	178
Gradients with respect to gates	179
Gradients with respect to weights	182
Gradients with respect to V	182
Gradients with respect to W	183
Gradients with respect to U	186
Predicting Bitcoin prices using LSTM model	187

1.1.1 Loading the data	Data preparation	188
1.1.2 Defining the parameters	Defining the parameters	191
1.1.3 Defining the LSTM cell	Define the LSTM cell	192
1.1.4 Defining forward propagation	Defining forward propagation	192
1.1.5 Defining backpropagation	Defining backpropagation	193
1.1.6 Training the LSTM model	Training the LSTM model	194
1.1.7 Making predictions using the LSTM model	Making predictions using the LSTM model	195
Gated recurrent units		
2.1.1 Understanding the GRU cell	Understanding the GRU cell	196
2.1.2 Update gate	Update gate	197
2.1.3 Reset gate	Reset gate	198
2.1.4 Updating hidden state	Updating hidden state	199
2.1.5 Forward propagation in a GRU cell	Forward propagation in a GRU cell	201
2.1.6 Backpropagation in a GRU cell	Backpropagation in a GRU cell	201
2.1.7 Gradient with respect to gates	Gradient with respect to gates	202
2.1.8 Gradients with respect to weights	Gradients with respect to weights	203
2.1.9 Gradients with respect to V	Gradients with respect to V	203
2.1.10 Gradients with respect to W	Gradients with respect to W	204
2.1.11 Gradients with respect to U	Gradients with respect to U	206
2.1.12 Implementing a GRU cell in TensorFlow	Implementing a GRU cell in TensorFlow	206
2.1.13 Defining the weights	Defining the weights	206
2.1.14 Defining forward propagation	Defining forward propagation	207
Bidirectional RNN		
Going deep with deep RNN		
Language translation using the seq2seq model		
3.1.1 Encoder	Encoder	212
3.1.2 Decoder	Decoder	214
3.1.3 Attention is all we need	Attention is all we need	217
Summary		
Questions		
Further reading		
Chapter 6: Demystifying Convolutional Networks		
What are CNNs?		
4.1.1 Convolutional layers	Convolutional layers	224
4.1.2 Strides	Strides	225
4.1.3 Padding	Padding	229
4.1.4 Pooling layers	Pooling layers	230
4.1.5 Fully connected layers	Fully connected layers	232
The architecture of CNNs		
The math behind CNNs		
5.1.1 Forward propagation	Forward propagation	234
5.1.2 Backward propagation	Backward propagation	235
Implementing a CNN in TensorFlow		
6.1.1 Defining helper functions	Defining helper functions	244
6.1.2 Defining the convolutional network	Defining the convolutional network	245
6.1.3 Computing loss	Computing loss	247

Starting the training	How to do a GAN	Introducing GANs	248
Visualizing extracted features		Generating images from words	248
CNN architectures		Building up to Word2Vec	250
LeNet architecture	Training CIFAR images	Loadable file dataset	251
Understanding AlexNet		Preprocessing and backpropagation	251
Architecture of VGGNet		Building the model	252
GoogleNet		Evaluating the embeddings	254
Inception v1		Visualizing word embeddings	255
Inception v2 and v3		Google	259
Capsule networks		Building a Capsule Network – Part 1	261
Understanding Capsule networks		Building a Capsule Network – Part 2	264
Computing prediction vectors		Finding similar documents using attributes	265
Coupling coefficients		Building a Capsule Network – Part 3	266
Squashing function		pick-thresholds for sentence embeddings	267
Dynamic routing algorithm		Summary	267
Architecture of the Capsule network		Further reading	268
The loss function			269
Margin loss			269
Reconstruction loss			270
Building Capsule networks in TensorFlow			271
Defining the squash function			271
Defining a dynamic routing algorithm			272
Computing primary and digit capsules			273
Masking the digit capsule			275
Defining the decoder			275
Computing the accuracy of the model			275
Calculating loss			276
Margin loss			276
Reconstruction loss			277
Total loss			277
Training the Capsule network			277
Summary			278
Questions			279
Further reading			279
Chapter 7: Learning Text Representations			
Understanding the word2vec model			
Understanding the CBOW model		Generating images with GAN	281
CBOW with a single context word	Forward propagation	Reading file dataset	282
	Backward propagation	Defining the generator	283
CBOW with multiple context words		Defining the discriminator	286
Understanding skip-gram model	Forward propagation in skip-gram	Defining the input layer	288
	Backward propagation	Sticking the GAN	291
Various training strategies		Comparing the loss functions	294
Hierarchical softmax		Discriminator loss	297
		Generator loss	298
		Obtaining the loss	302
		Getting the gradients	303
		Generating new words	303

8.1	Negative sampling	305
8.2	Subsampling frequent words	306
Building the word2vec model using gensim		306
8.3	Loading the dataset	307
8.4	Preprocessing and preparing the dataset	308
8.5	Building the model	310
8.6	Evaluating the embeddings	311
Visualizing word embeddings in TensorBoard		312
Doc2vec		315
8.7	Paragraph Vector – Distributed Memory model	315
8.8	Paragraph Vector – Distributed Bag of Words model	316
8.9	Finding similar documents using doc2vec	317
Understanding skip-thoughts algorithm		320
Quick-thoughts for sentence embeddings		321
Summary		322
Questions		323
Further reading		323

Section 3: Advanced Deep Learning Algorithms

Chapter 8: Generating Images Using GANs		327
Differences between discriminative and generative models		328
Say hello to GANs!		329
8.1	Breaking down the generator	331
8.2	Breaking down the discriminator	332
8.3	How do they learn though?	333
8.4	Architecture of a GAN	335
8.5	Demystifying the loss function	335
8.6	Discriminator loss	336
8.7	First term	336
8.8	Second term	337
8.9	Final term	338
8.10	Generator loss	338
8.11	Total loss	338
8.12	Heuristic loss	339
8.13	Generating images using GANs in TensorFlow	340
8.14	Reading the dataset	341
8.15	Defining the generator	341
8.16	Defining the discriminator	342
8.17	Defining the input placeholders	342
8.18	Starting the GAN!	342
8.19	Computing the loss function	342
8.20	Discriminator loss	343
8.21	Generator loss	344
8.22	Optimizing the loss	345
8.23	Starting the training	345
8.24	Generating handwritten digits	345

DCGAN – Adding convolution to a GAN	348
Deconvolutional generator	348
Convolutional discriminator	349
Implementing DCGAN to generate CIFAR images	350
Exploring the dataset	351
Defining the discriminator	352
Defining the generator	353
Defining the inputs	355
Starting the DCGAN	355
Computing the loss function	355
Discriminator loss	356
Generator loss	356
Optimizing the loss	356
Train the DCGAN	357
Least squares GAN	359
Loss function	361
LSGAN in TensorFlow	362
Discriminator loss	362
Generator loss	362
GANs with Wasserstein distance	363
Are we minimizing JS divergence in GANs?	363
What is the Wasserstein distance?	366
Demystifying the k-Lipschitz function	367
The loss function of WGAN	368
WGAN in TensorFlow	369
Summary	370
Questions	371
Further reading	371
Chapter 9: Learning More about GANs	373
Conditional GANs	374
Loss function of CGAN	376
Generating specific handwritten digits using CGAN	376
Defining the generator	377
Defining discriminator	378
Start the GAN!	378
Computing the loss function	379
Discriminator loss	379
Generator loss	380
Optimizing the loss	380
Start training the CGAN	380
Generate the handwritten digit, 7	381
Understanding InfoGAN	382
Mutual information	383
Architecture of the InfoGAN	385
Constructing an InfoGAN in TensorFlow	386
Defining generator	387
Defining the discriminator	388

Define the input placeholders	390
Start the GAN	390
Computing loss function	391
Discriminator loss	391
Generator loss	391
Mutual information	392
Optimizing the loss	393
Beginning training	393
Generating handwritten digits	394
Translating images using a CycleGAN	395
Role of generators	398
Role of discriminators	399
Loss function	400
Cycle consistency loss	401
Converting photos to paintings using a CycleGAN	403
StackGAN	407
The architecture of StackGANs	408
Conditioning augmentation	409
Stage I	409
Generator	409
Discriminator	410
Stage II	410
Generator	411
Discriminator	411
Summary	411
Questions	412
Further reading	412
Chapter 10: Reconstructing Inputs Using Autoencoders	413
What is an autoencoder?	414
Understanding the architecture of autoencoders	417
Reconstructing the MNIST images using an autoencoder	418
Preparing the dataset	419
Defining the encoder	419
Defining the decoder	420
Building the model	420
Reconstructing images	421
Plotting reconstructed images	421
Autoencoders with convolutions	422
Building a convolutional autoencoder	423
Defining the encoder	424
Defining the decoder	425
Building the model	425
Reconstructing the images	426
Exploring denoising autoencoders	427
Denoising images using DAE	428
Understanding sparse autoencoders	430
Building the sparse autoencoder	433

Defining the sparse regularizer	433
Learning to use contractive autoencoders	434
Implementing the contractive autoencoder	434
Defining the contractive loss	435
Dissecting variational autoencoders	436
Variational inference	438
The loss function	439
Reparameterization trick	441
Generating images using VAE	442
Preparing the dataset	443
Defining the encoder	443
Defining the sampling operation	444
Defining the decoder	444
Building the model	444
Defining the generator	445
Plotting generated images	445
Summary	446
Questions	447
Further reading	447
Chapter 11: Exploring Few-Shot Learning Algorithms	449
What is few-shot learning?	450
Siamese networks	451
Architecture of siamese networks	452
Prototypical networks	454
Relation networks	460
Matching networks	462
Support set embedding function	465
Query set embedding function	465
The architecture of matching networks	466
Summary	467
Questions	468
Further reading	468
Appendix A: Assessments	469
Chapter 1 - Introduction to Deep Learning	469
Chapter 2 - Getting to Know TensorFlow	470
Chapter 3 - Gradient Descent and Its Variants	471
Chapter 4 - Generating Song Lyrics Using an RNN	472
Chapter 5 - Improvements to the RNN	473
Chapter 6 - Demystifying Convolutional Networks	474
Chapter 7 - Learning Text Representations	475
Chapter 8 - Generating Images Using GANs	476
Chapter 9 - Learning More about GANs	477
Chapter 10 - Reconstructing Inputs Using Autoencoders	478

Table of Contents