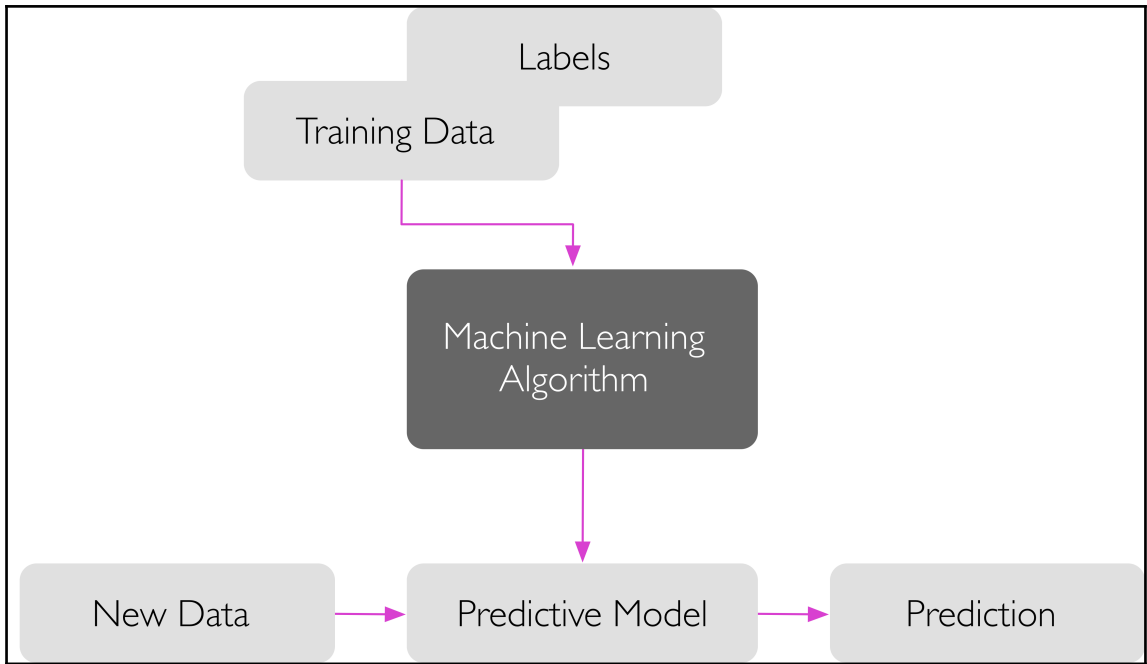

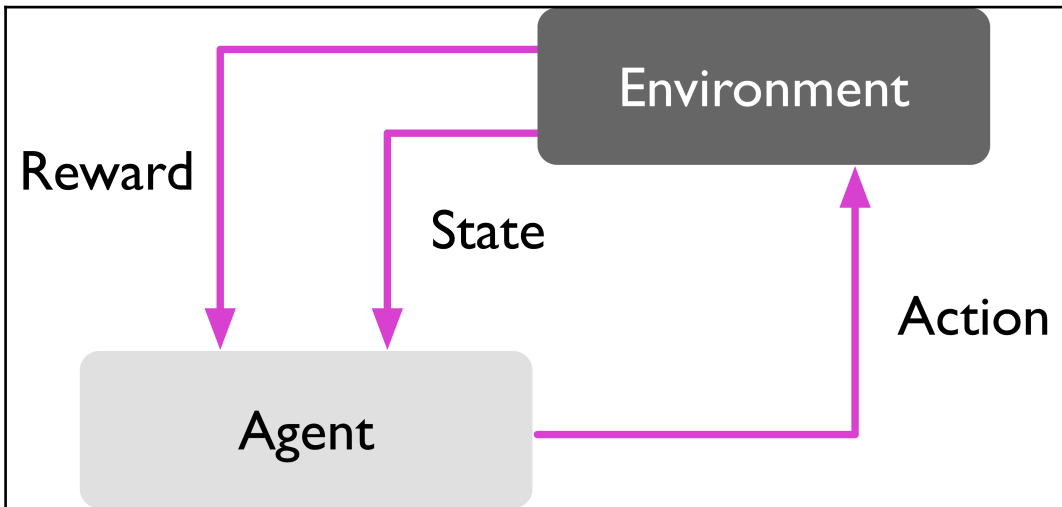





Chapter 1: Giving Computers the Ability to Learn from Data


Supervised Learning	<ul style="list-style-type: none">> Labeled data> Direct feedback> Predict outcome/future
Unsupervised Learning	<ul style="list-style-type: none">> No labels> No feedback> Find hidden structure in data
Reinforcement Learning	<ul style="list-style-type: none">> Decision process> Reward system> Learn series of actions

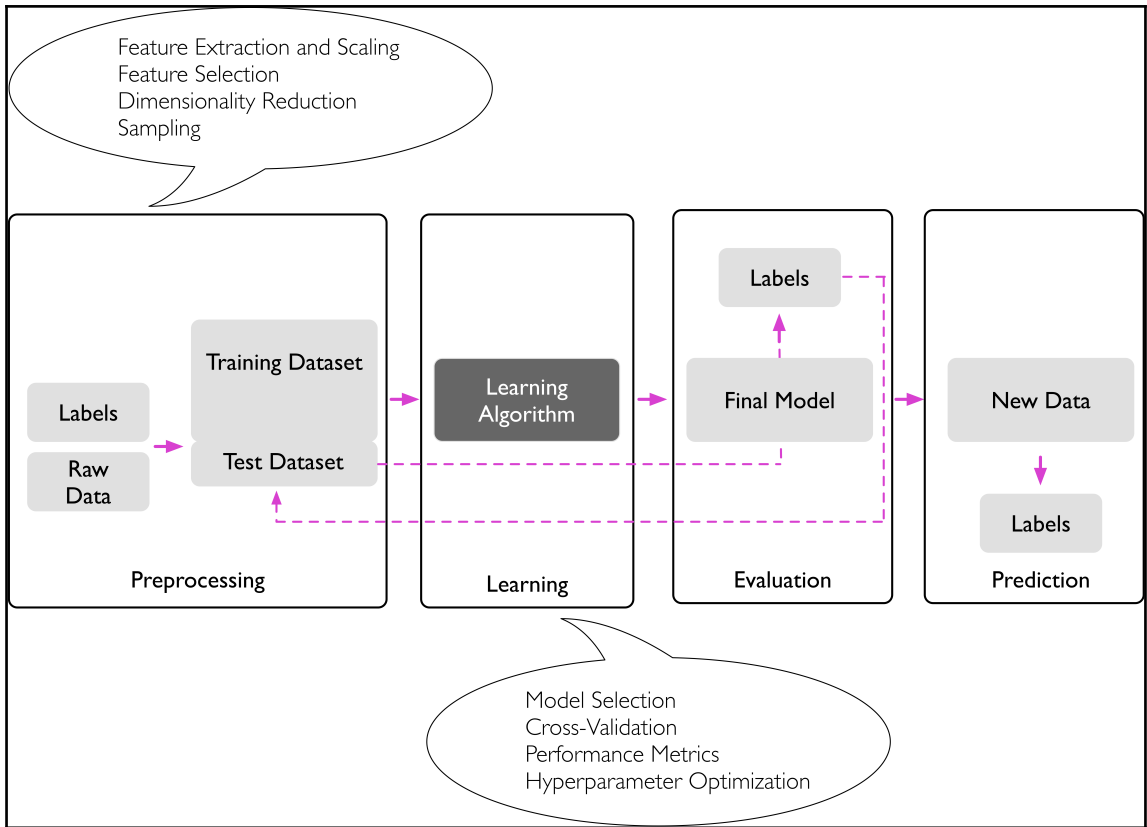








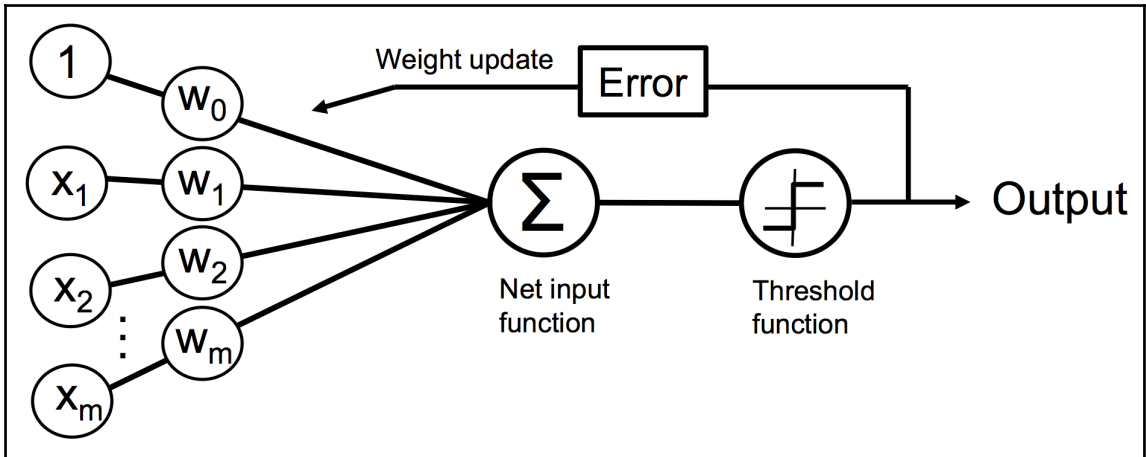





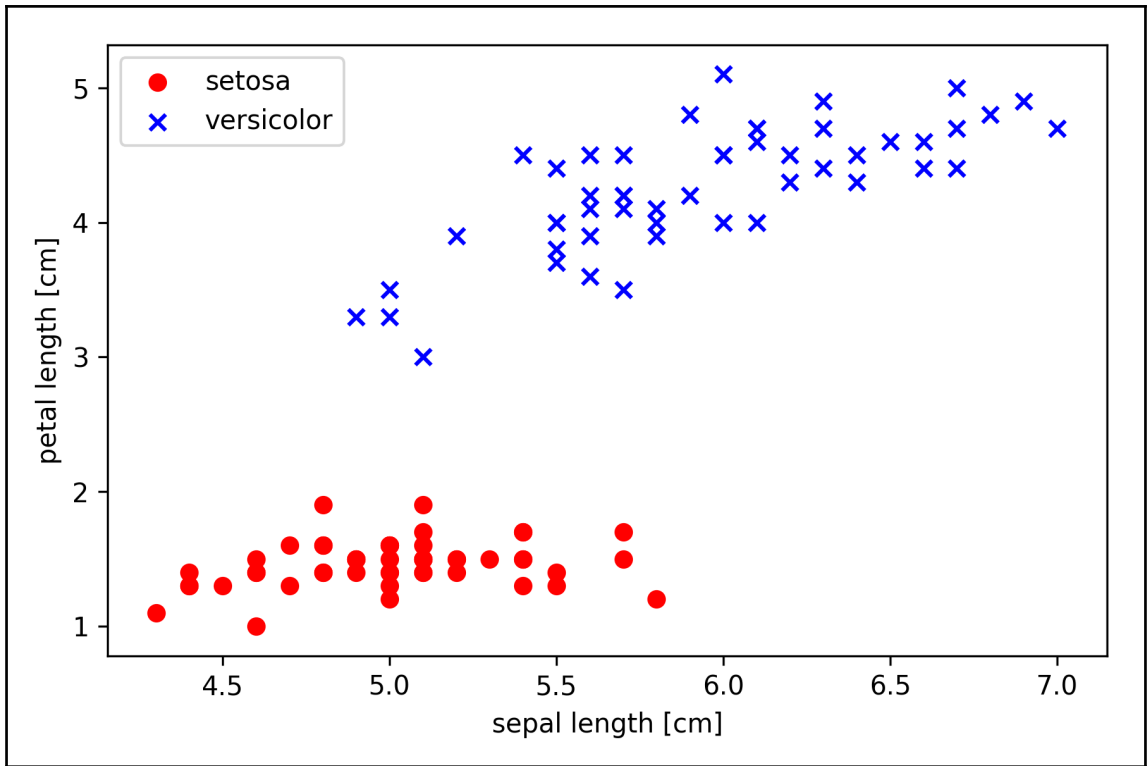


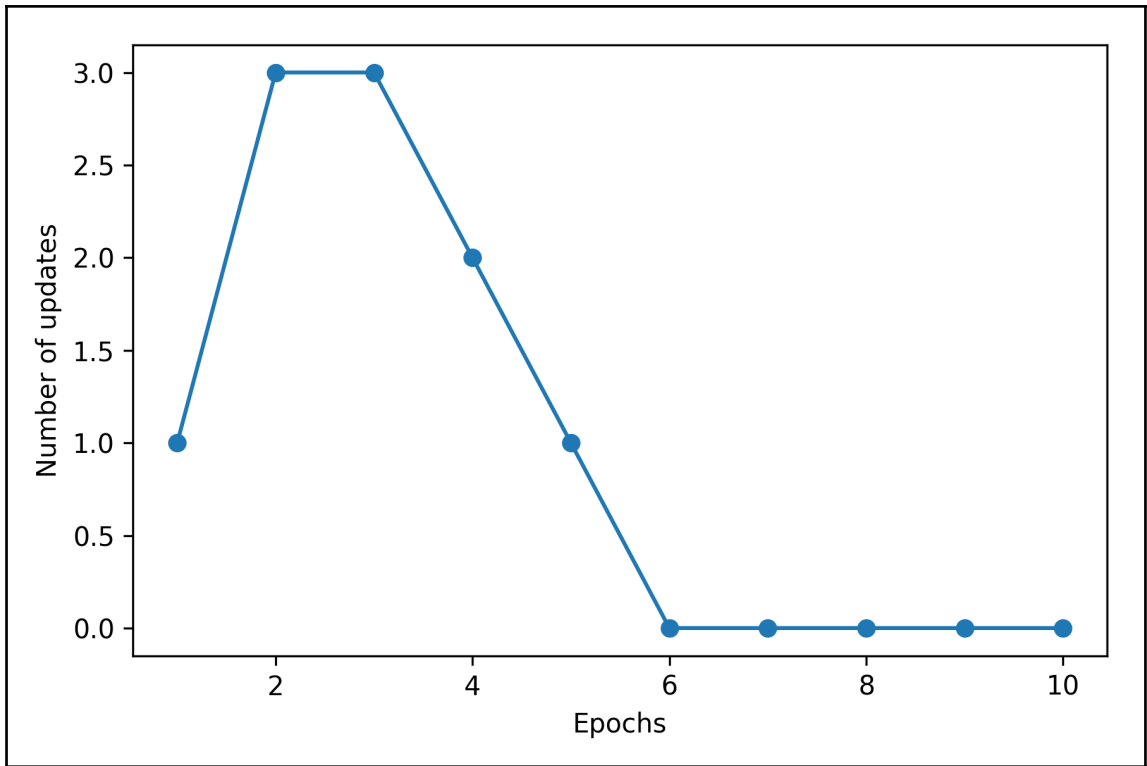
Chapter 2: Training Simple Machine Learning Algorithms for Classification

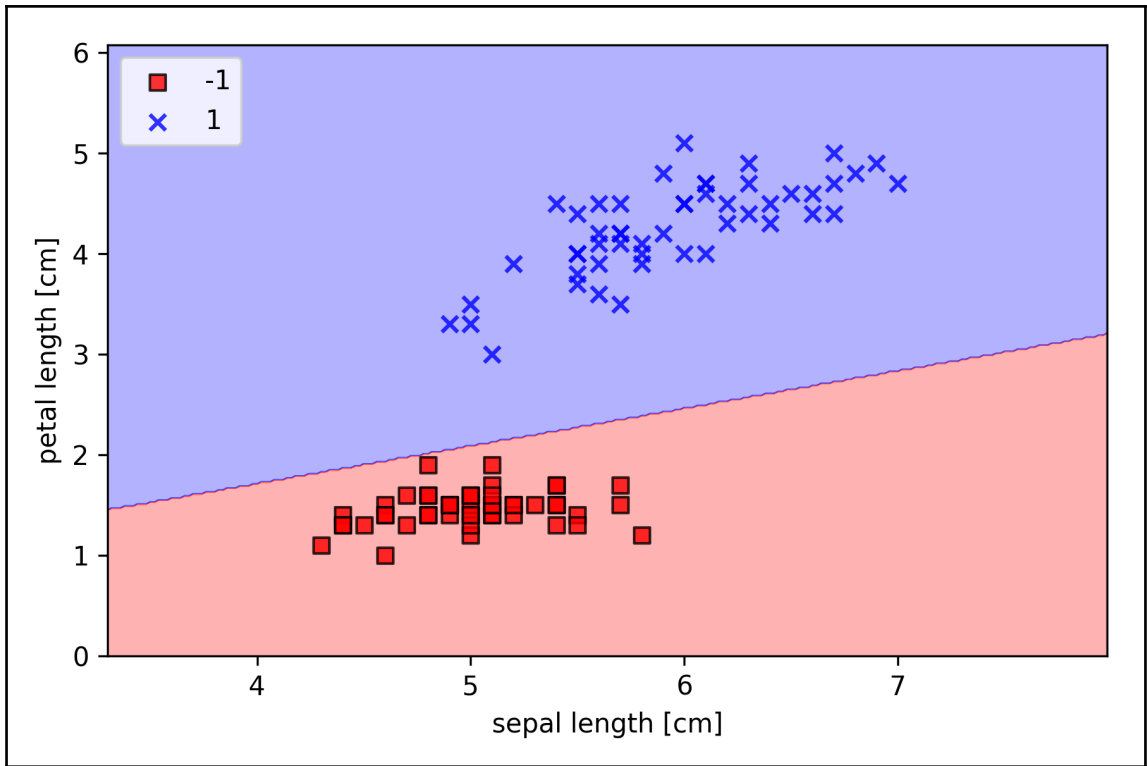


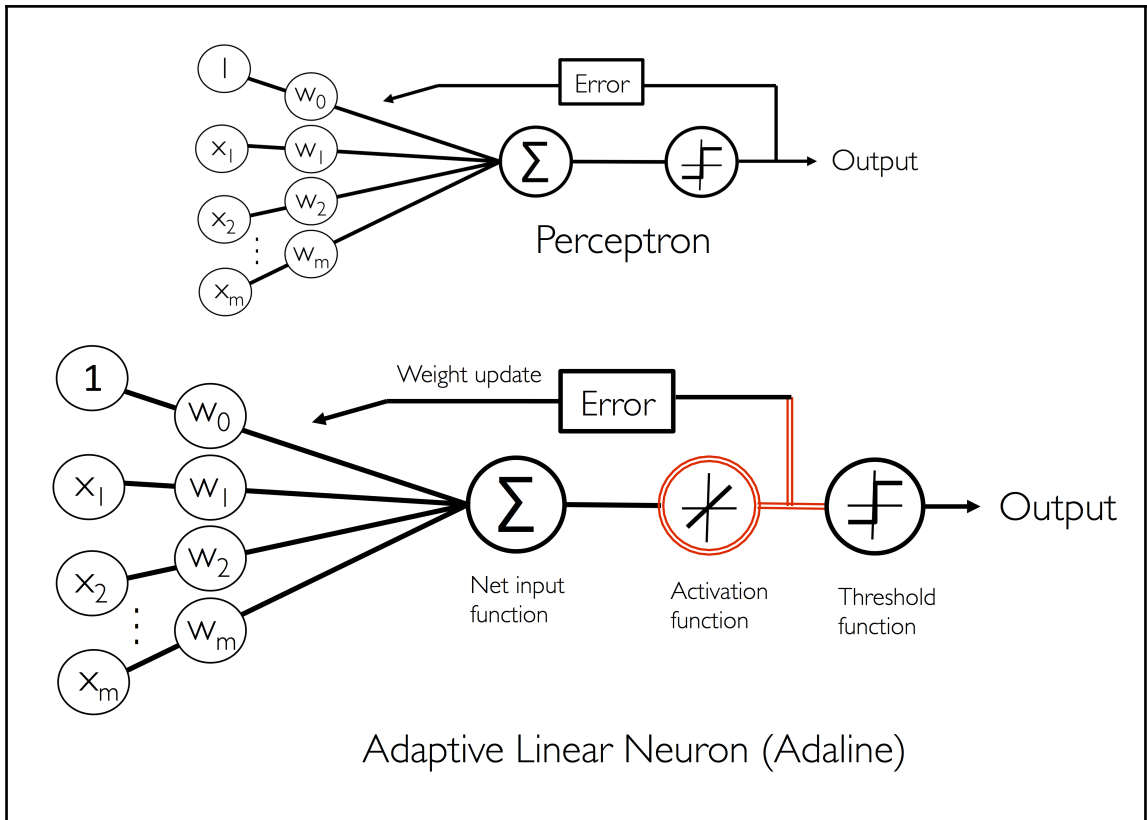


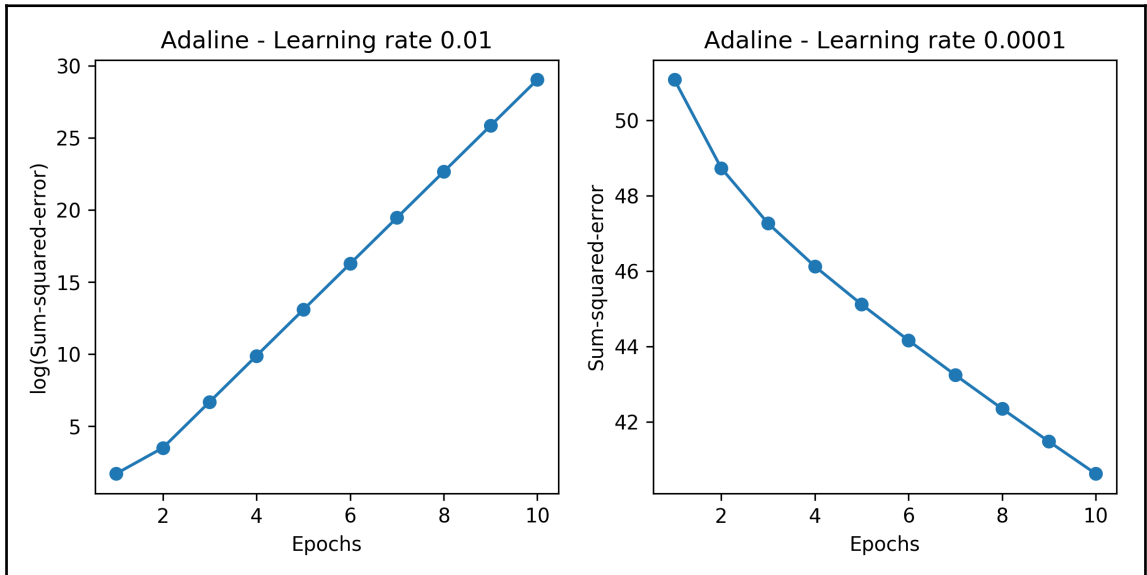

	0	1	2	3	4
145	6.7	3.0	5.2	2.3	Iris-virginica
146	6.3	2.5	5.0	1.9	Iris-virginica
147	6.5	3.0	5.2	2.0	Iris-virginica
148	6.2	3.4	5.4	2.3	Iris-virginica
149	5.9	3.0	5.1	1.8	Iris-virginica

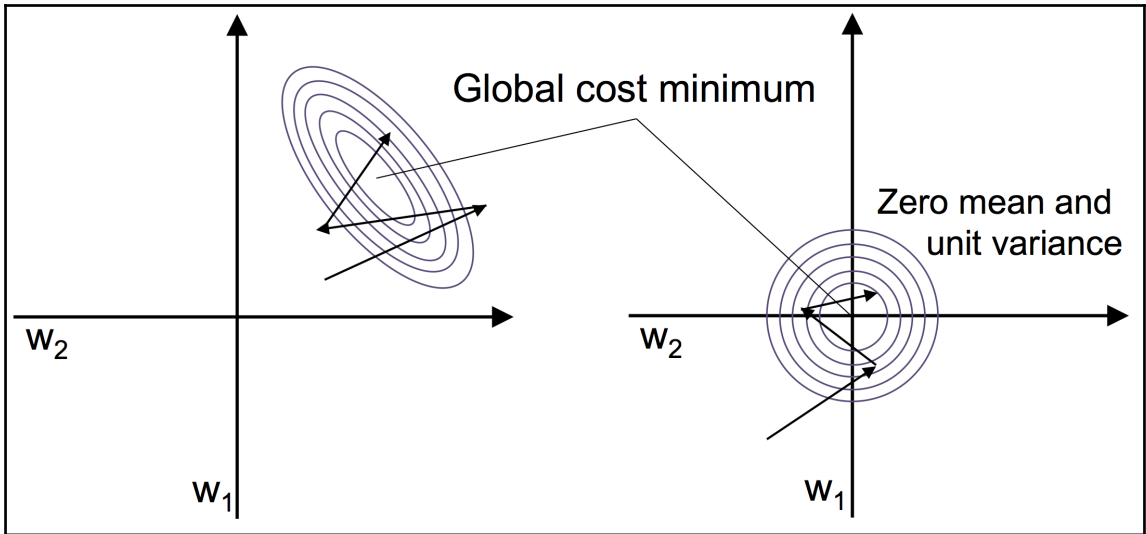




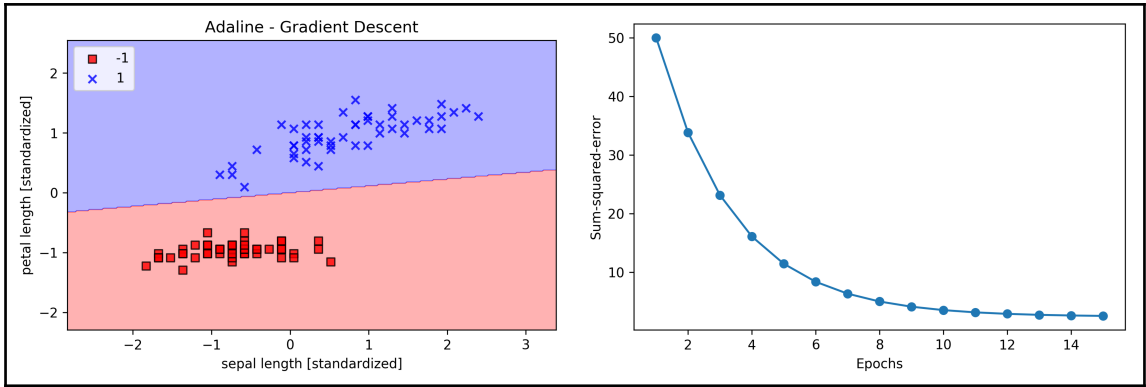




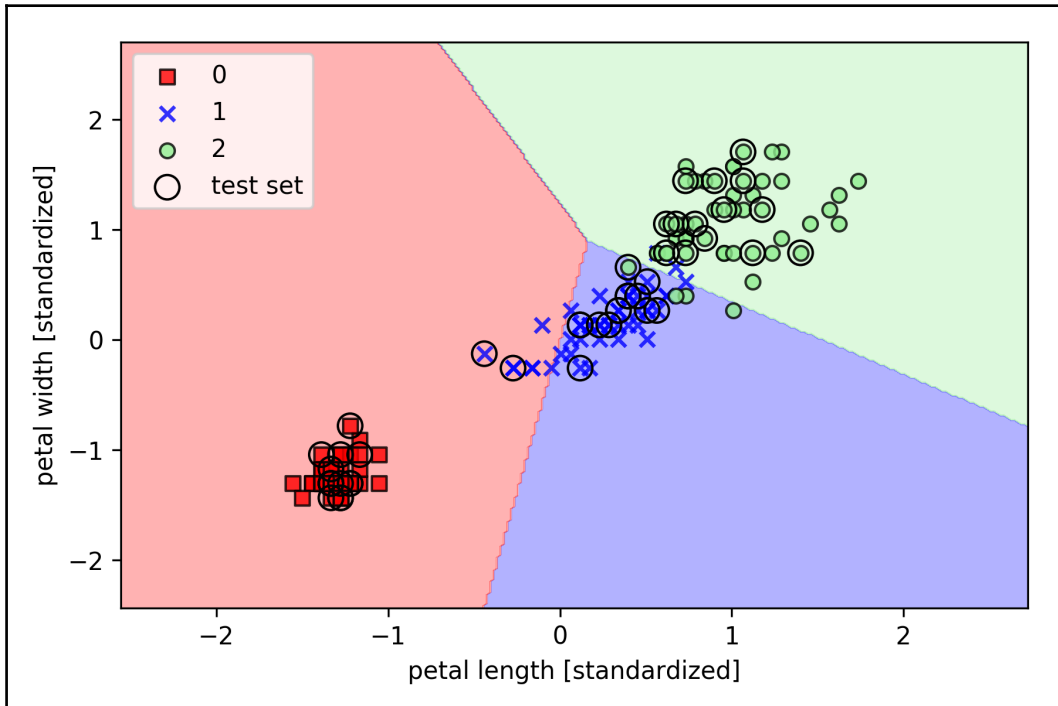





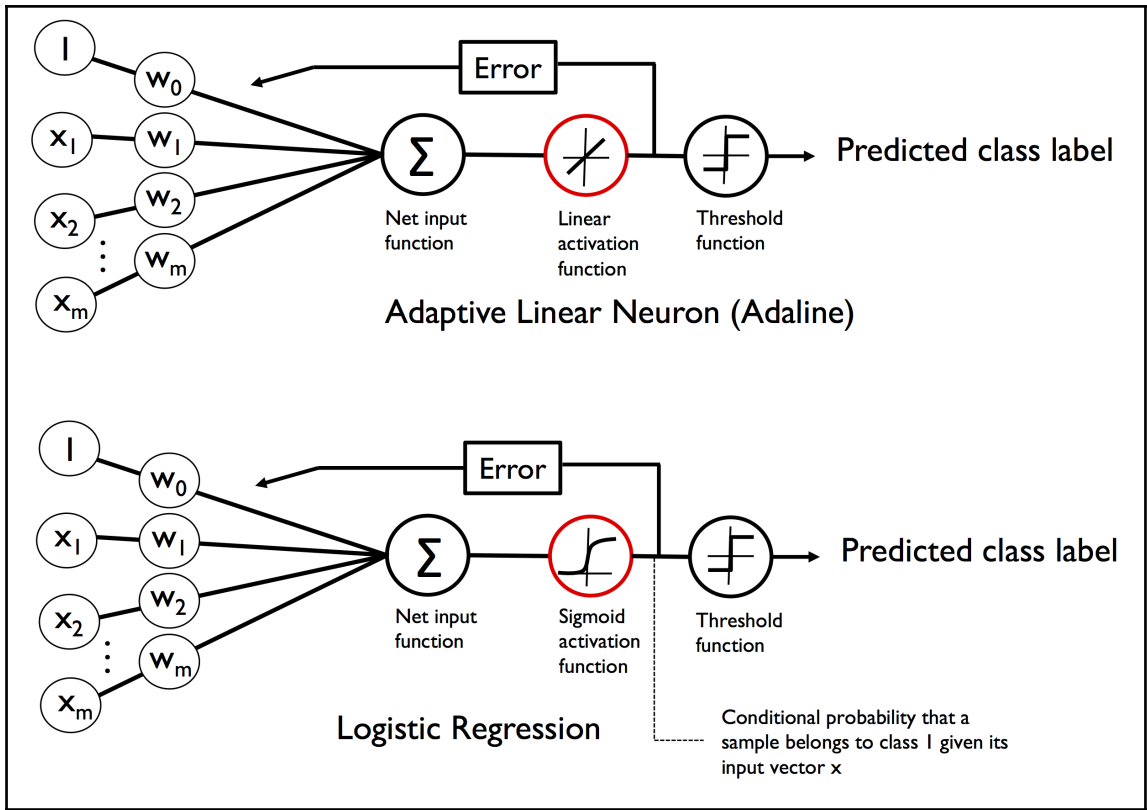






Chapter 3: A Tour of Machine Learning Classifiers Using scikit-learn



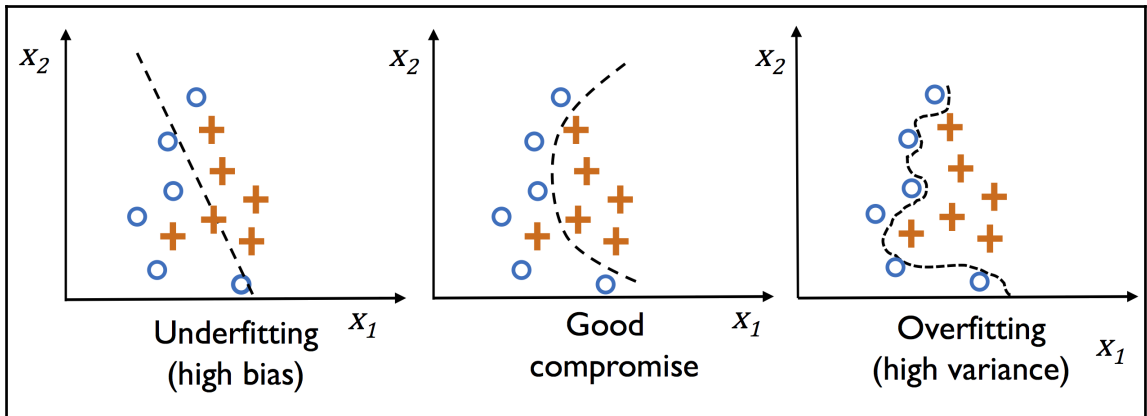



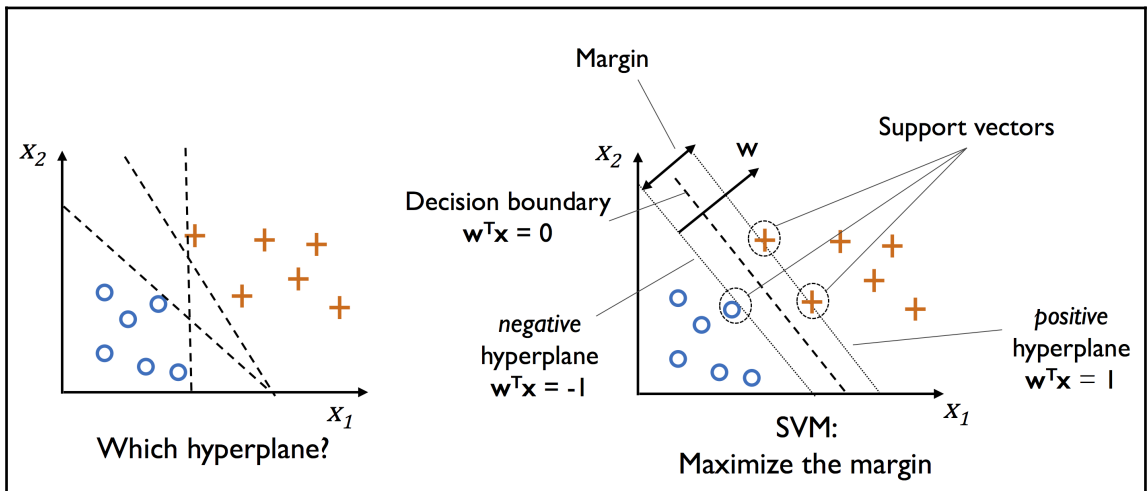
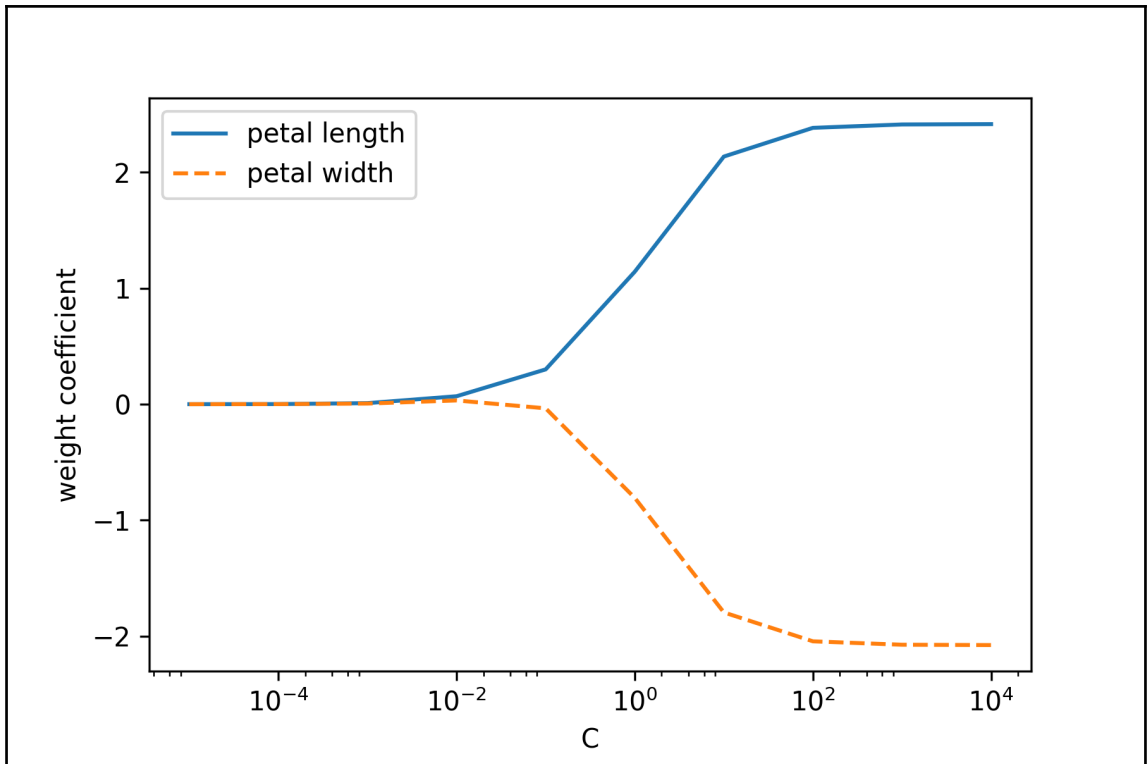









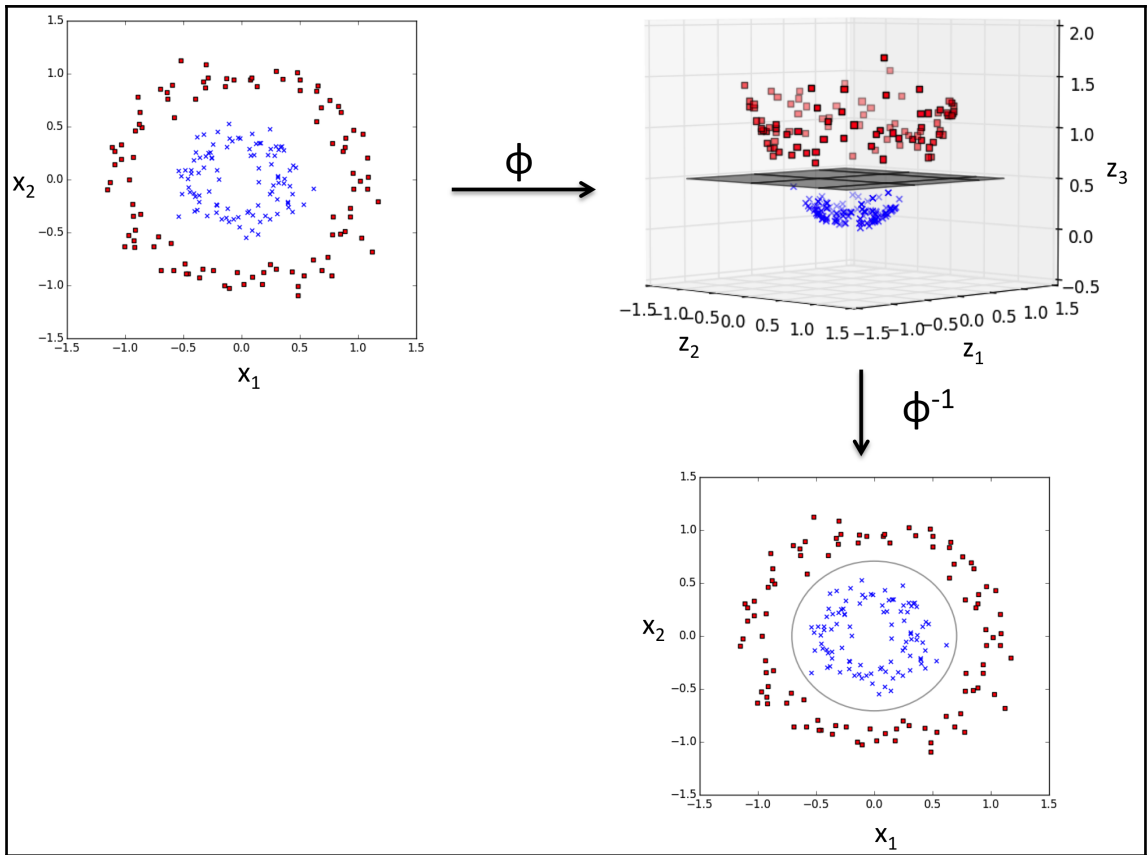
>

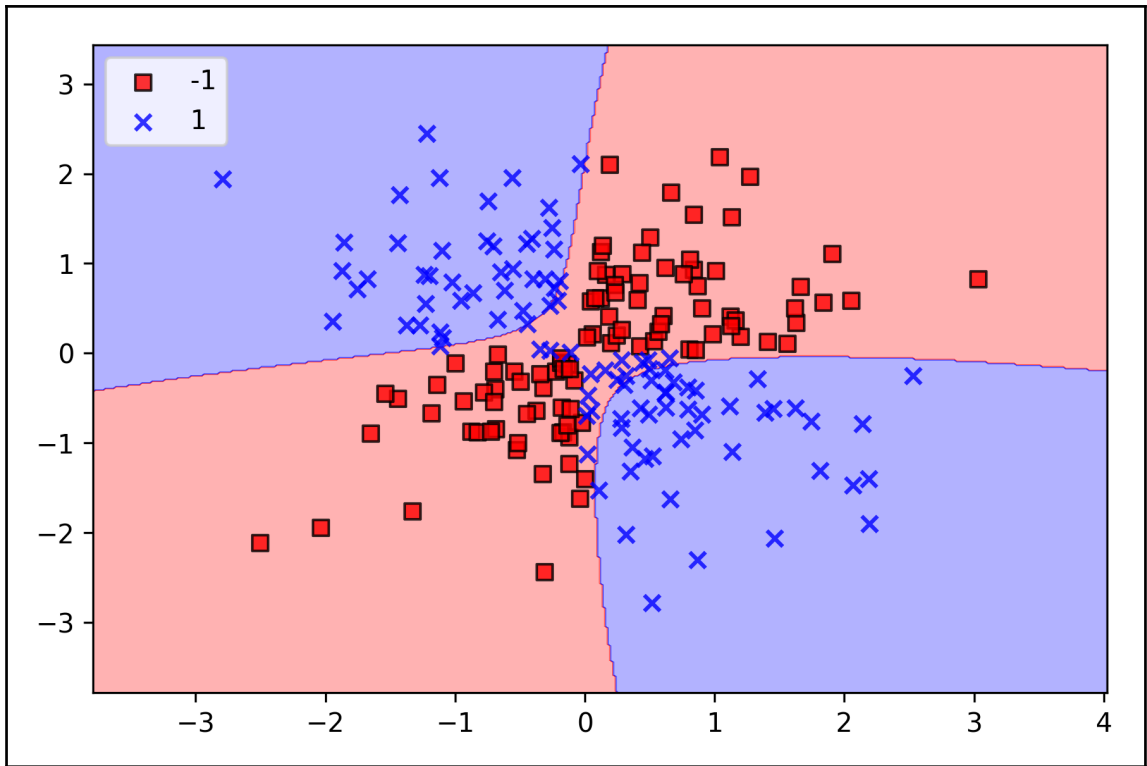









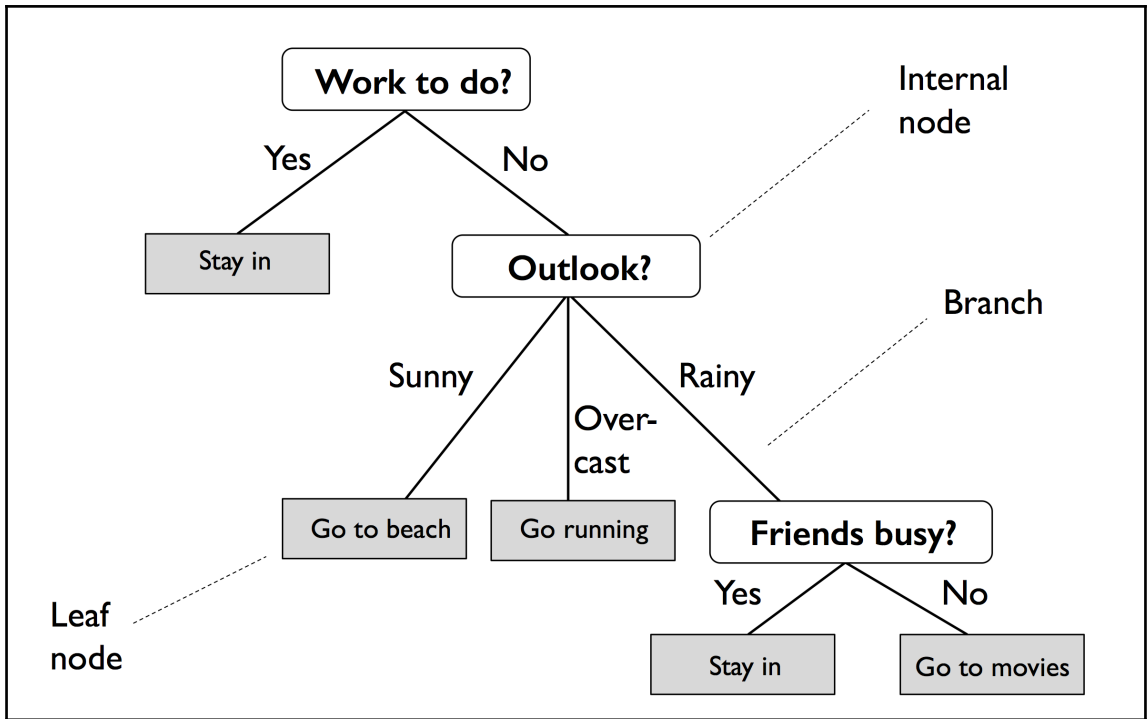








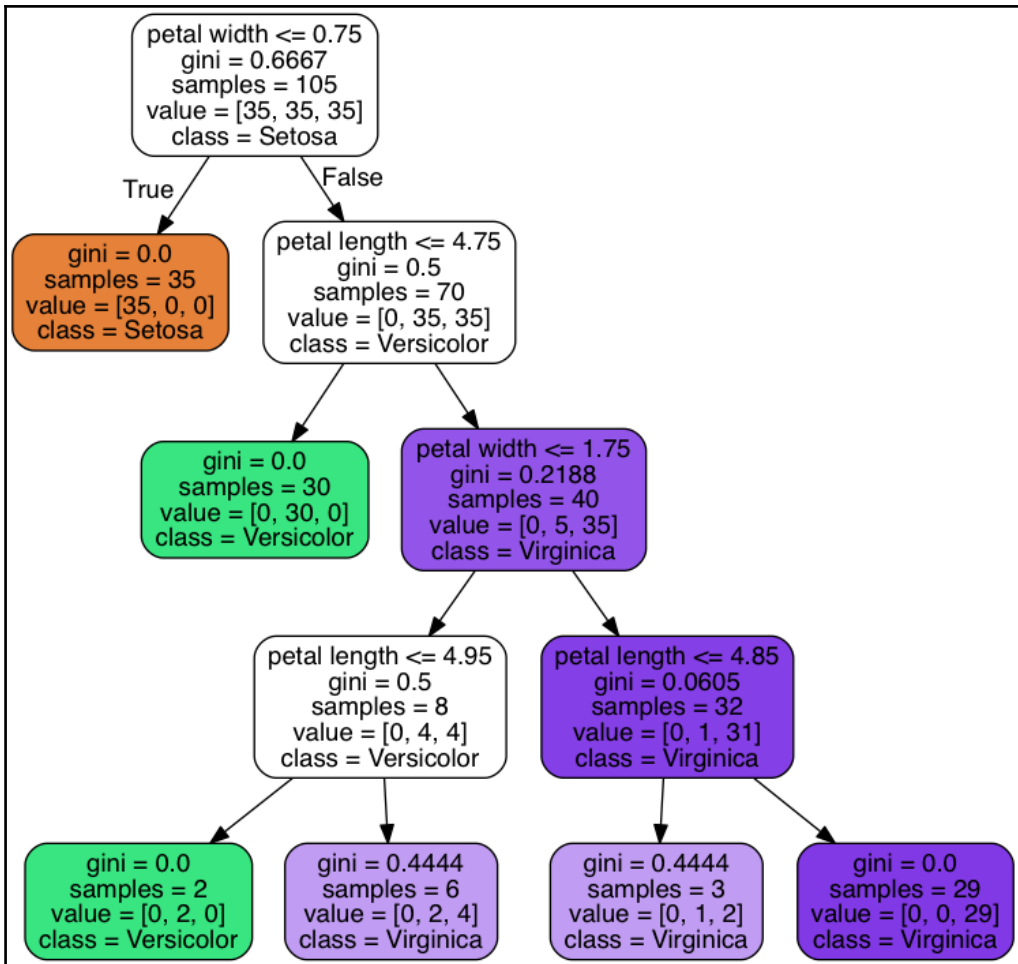













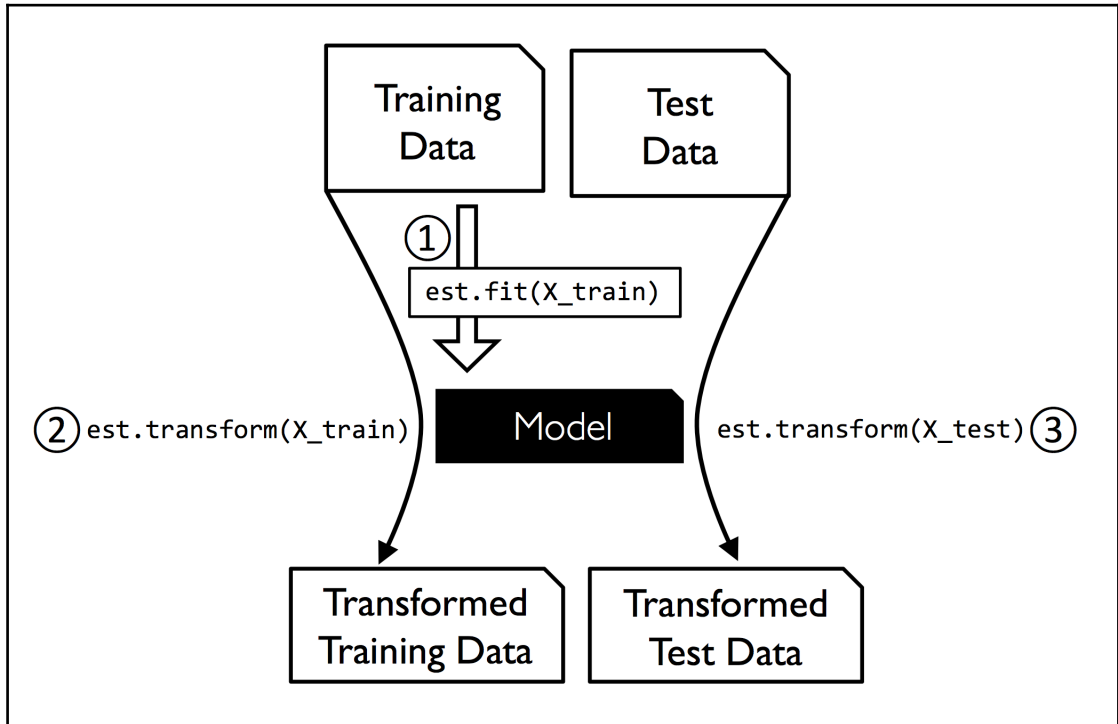


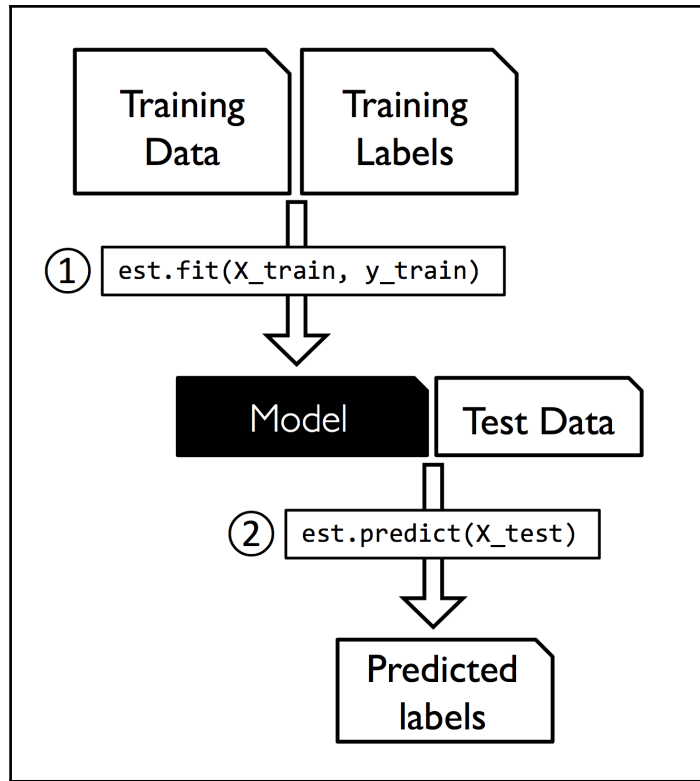








Chapter 4: Building Good Training Sets – Data Preprocessing









	Class label	Alcohol	Malic acid	Ash	Alcalinity of ash	Magnesium	Total phenols	Flavanoids	Nonflavanoid phenols	Proanthocyanins	Color intensity	Hue	OD280/OD315 of diluted wines	Proline
0	1	14.23	1.71	2.43	15.6	127	2.80	3.06	0.28	2.29	5.64	1.04	3.92	1065
1	1	13.20	1.78	2.14	11.2	100	2.65	2.76	0.26	1.28	4.38	1.05	3.40	1050
2	1	13.16	2.36	2.67	18.6	101	2.80	3.24	0.30	2.81	5.68	1.03	3.17	1185
3	1	14.37	1.95	2.50	16.8	113	3.85	3.49	0.24	2.18	7.80	0.86	3.45	1480
4	1	13.24	2.59	2.87	21.0	118	2.80	2.69	0.39	1.82	4.32	1.04	2.93	735



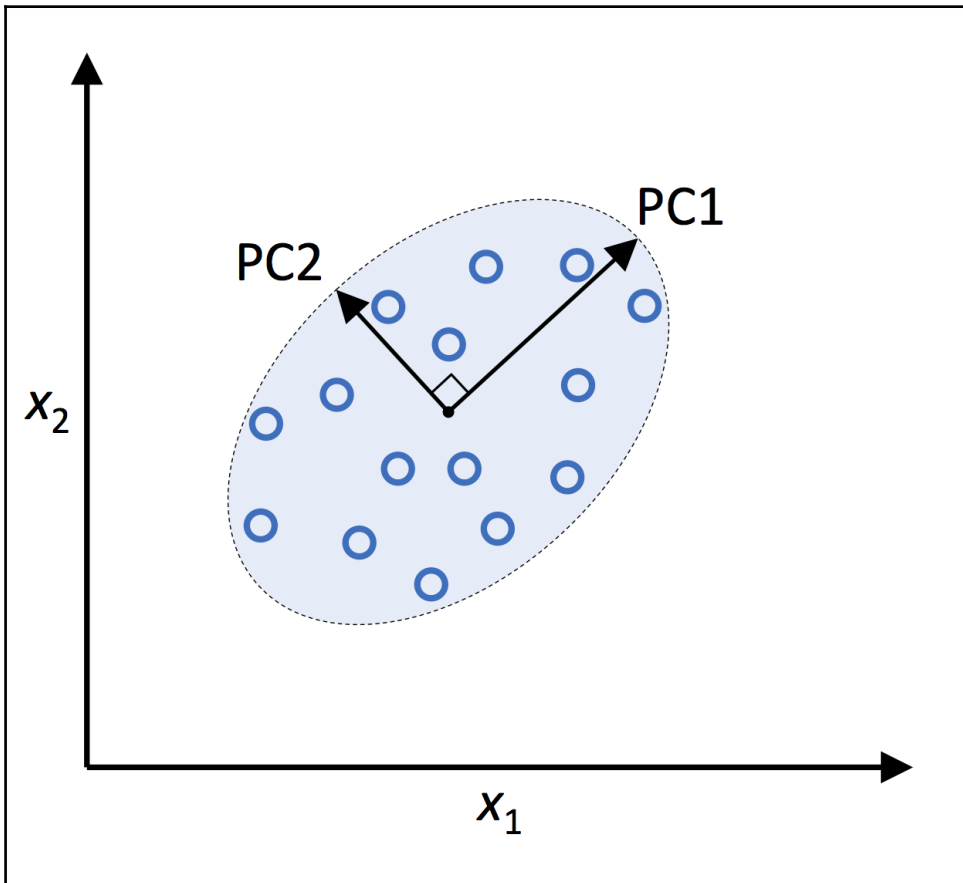












Chapter 5: Compressing Data via Dimensionality Reduction





































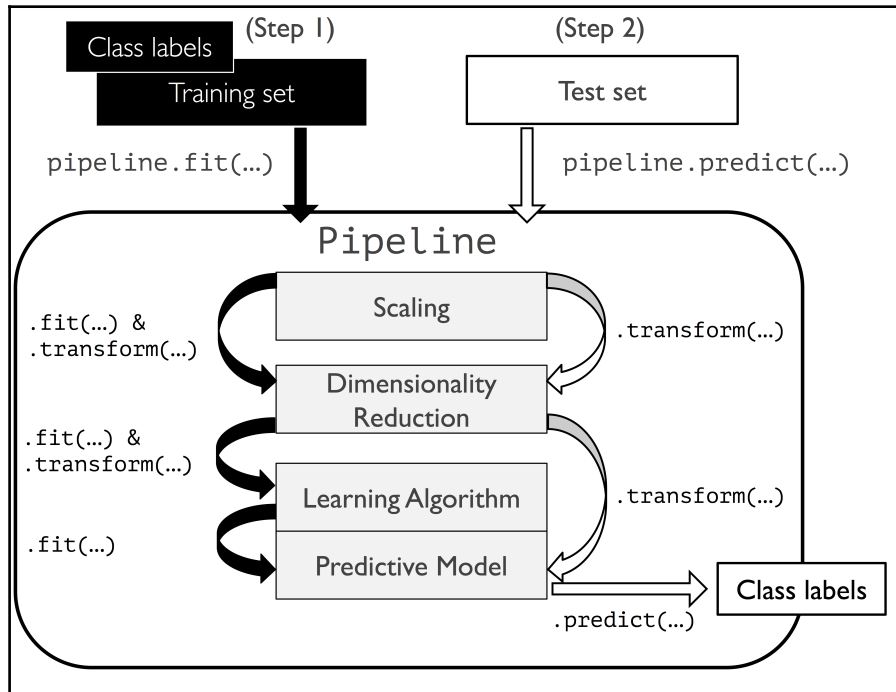


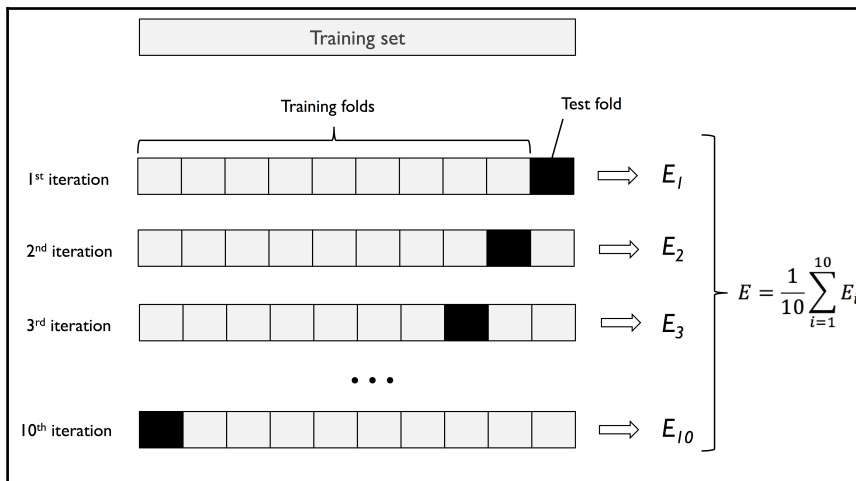
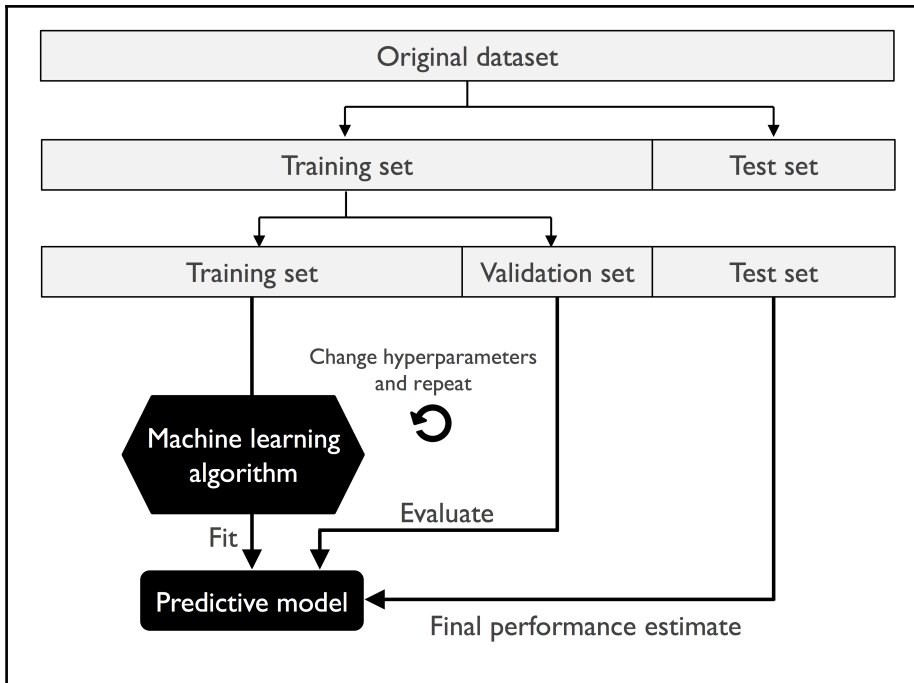


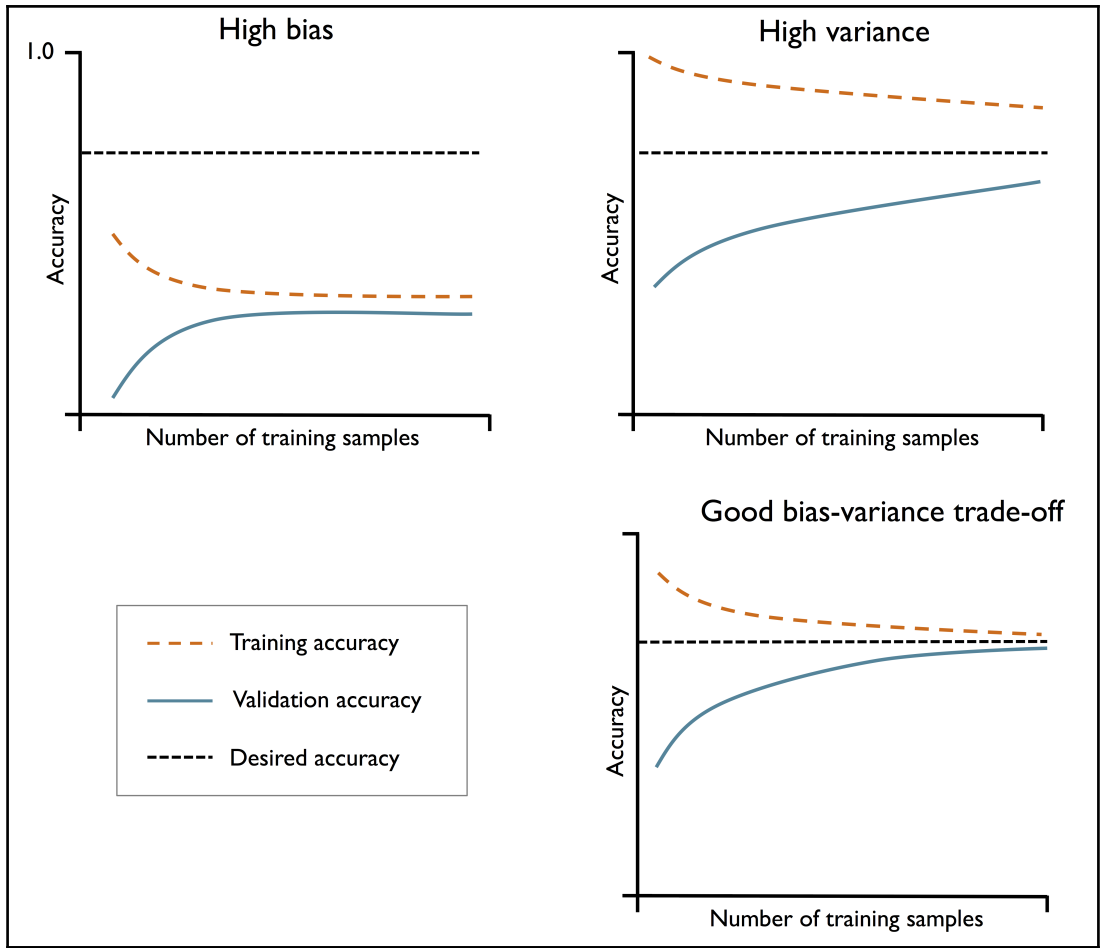


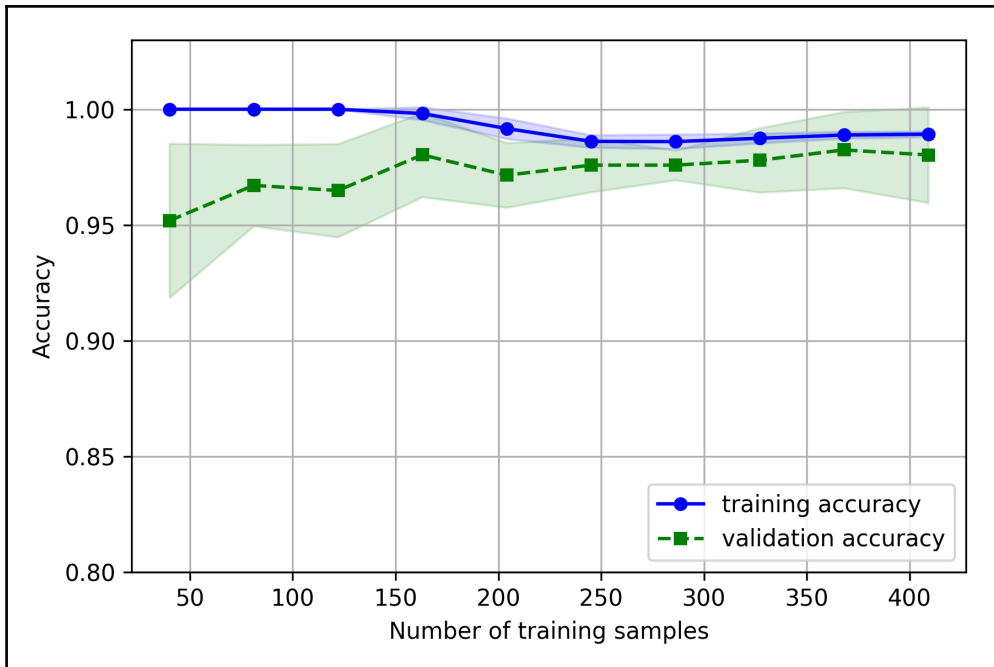


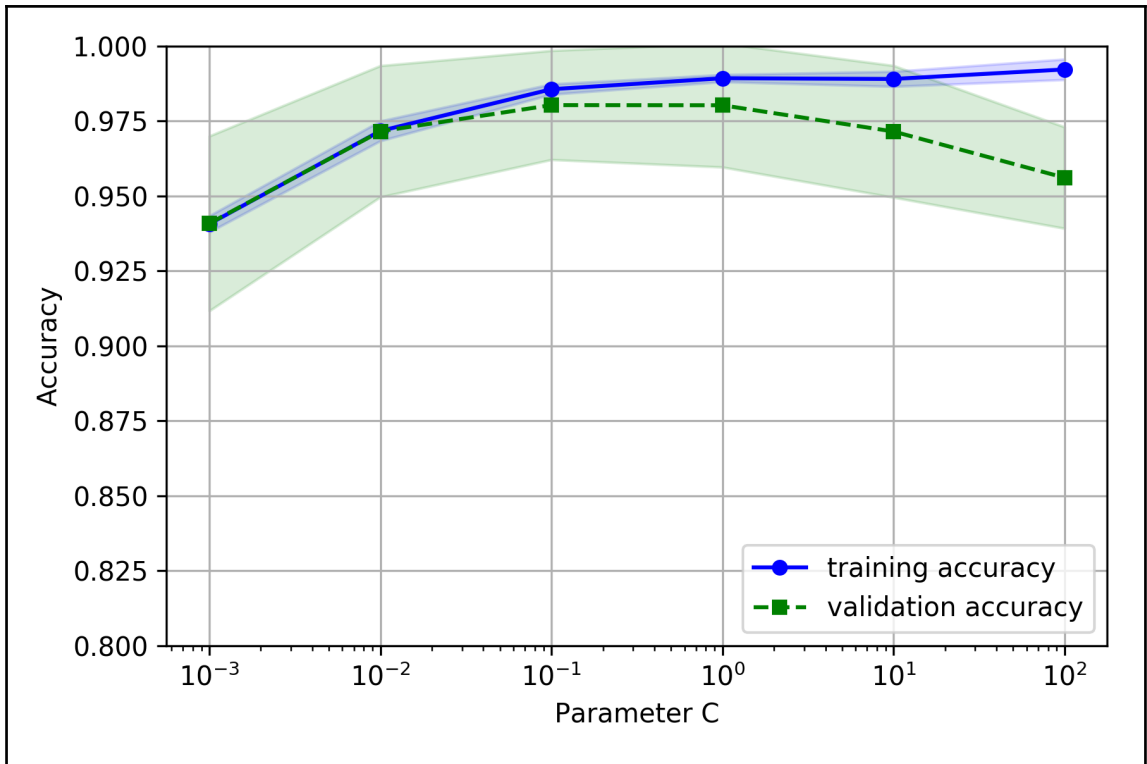
Chapter 6: Learning Best Practices for Model Evaluation and Hyperparameter Tuning

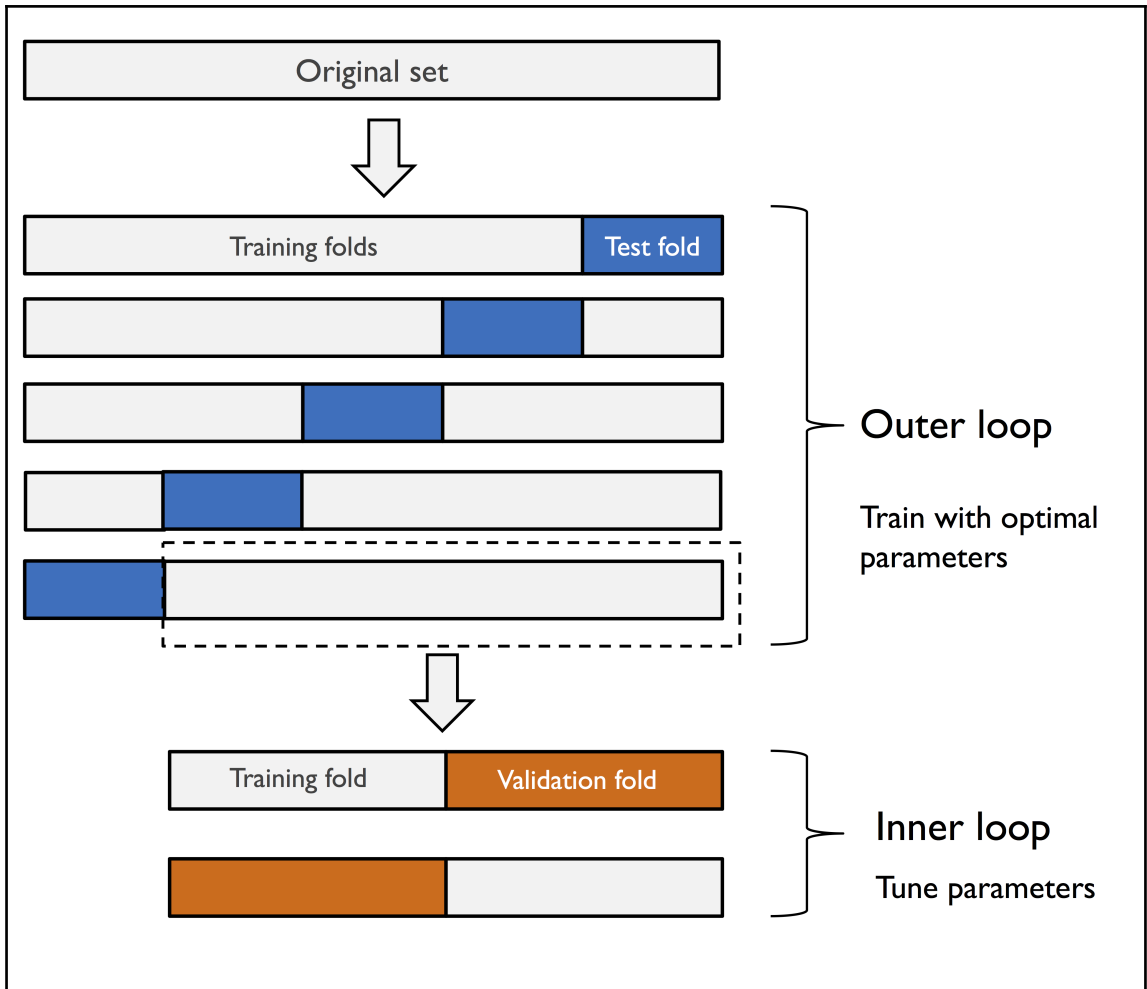












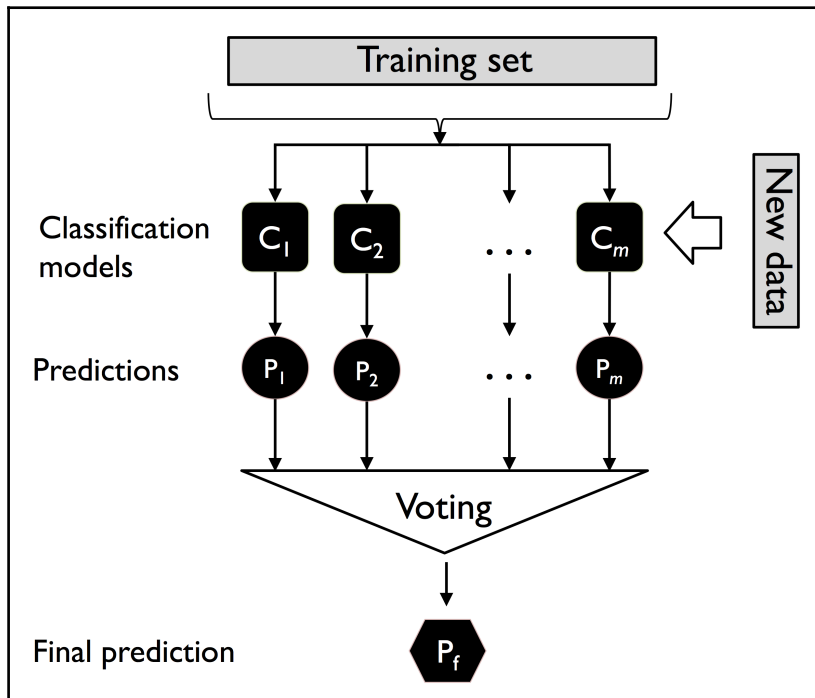




		Predicted class	
		<i>P</i>	<i>N</i>
Actual class	<i>P</i>	True positives (TP)	False negatives (FN)
	<i>N</i>	False positives (FP)	True negatives (TN)







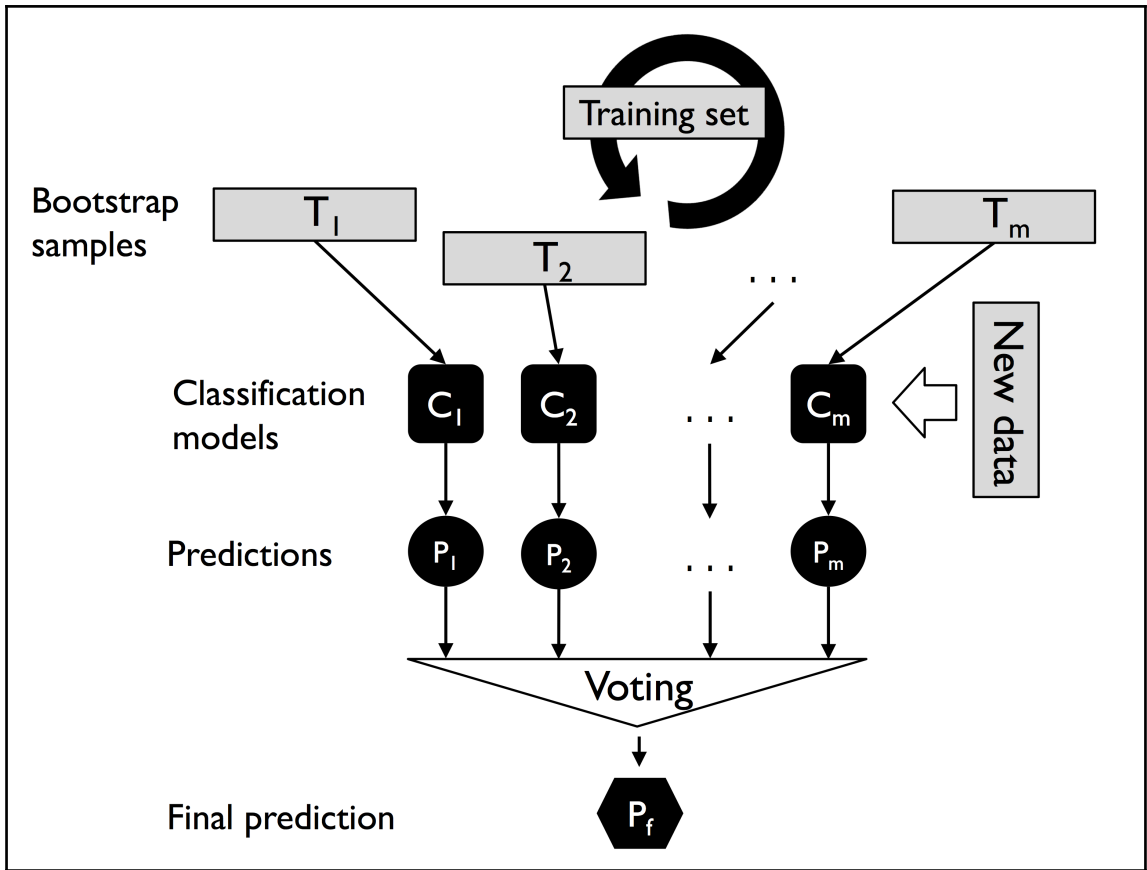
Chapter 7: Combining Different Models for Ensemble Learning






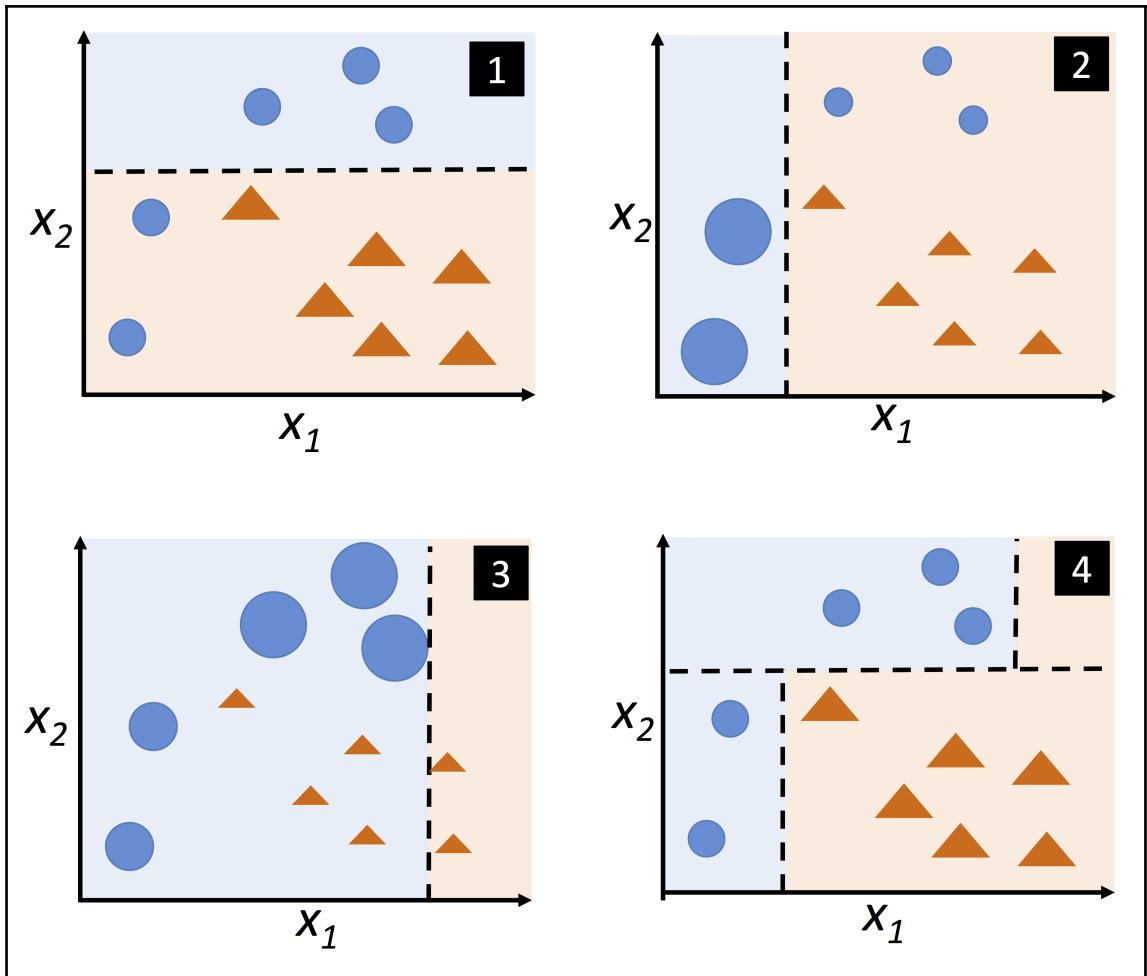







Sample indices	Bagging round 1	Bagging round 2	...
1	2	7	...
2	2	3	...
3	1	2	...
4	3	1	...
5	7	1	...
6	2	7	...
7	4	7	...





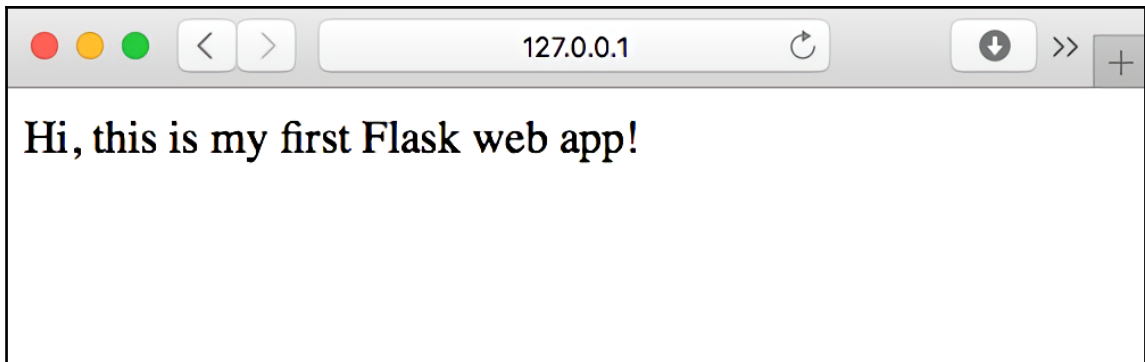
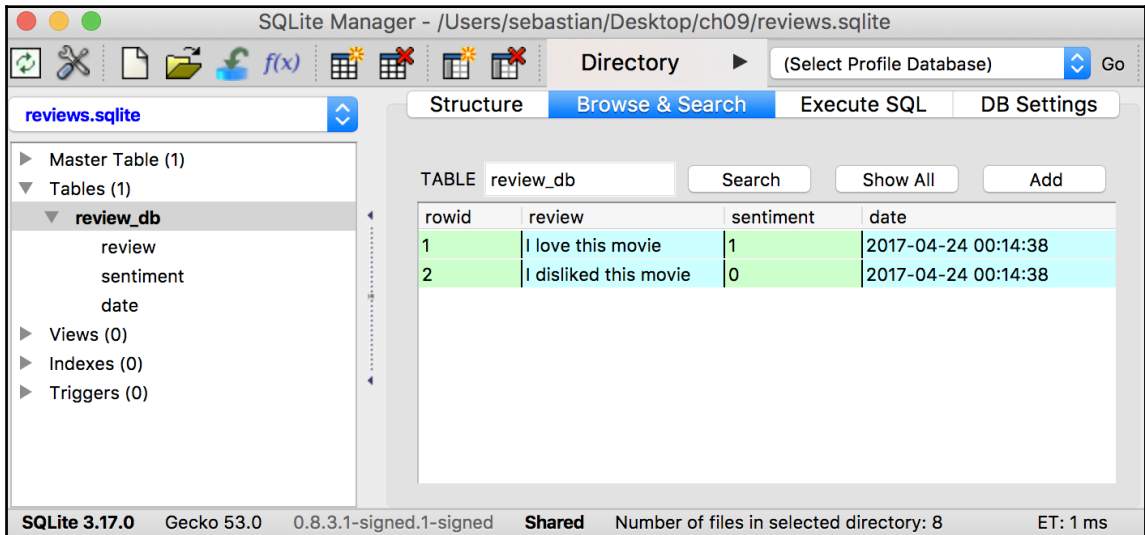
Sample indices	x	y	Weights	$\hat{y}(x \leq 3.0)?$	Correct?	Updated weights
1	1.0	1	0.1	1	Yes	0.072
2	2.0	1	0.1	1	Yes	0.072
3	3.0	1	0.1	1	Yes	0.072
4	4.0	-1	0.1	-1	Yes	0.072
5	5.0	-1	0.1	-1	Yes	0.072
6	6.0	-1	0.1	-1	Yes	0.072
7	7.0	1	0.1	-1	No	0.167
8	8.0	1	0.1	-1	No	0.167
9	9.0	1	0.1	-1	No	0.167
10	10.0	-1	0.1	-1	Yes	0.072

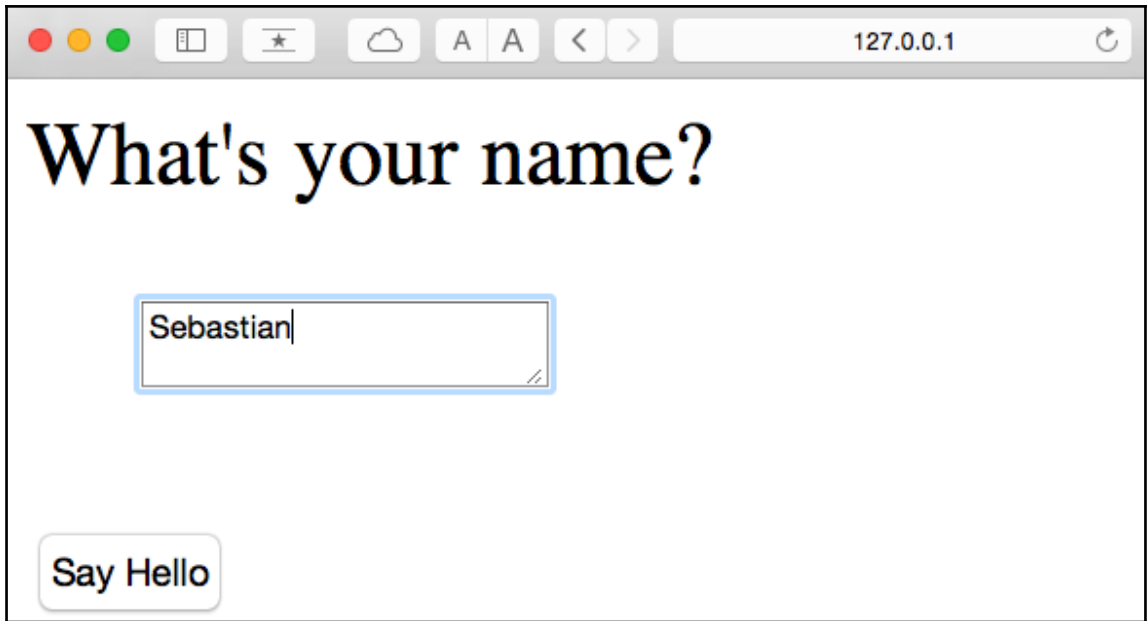


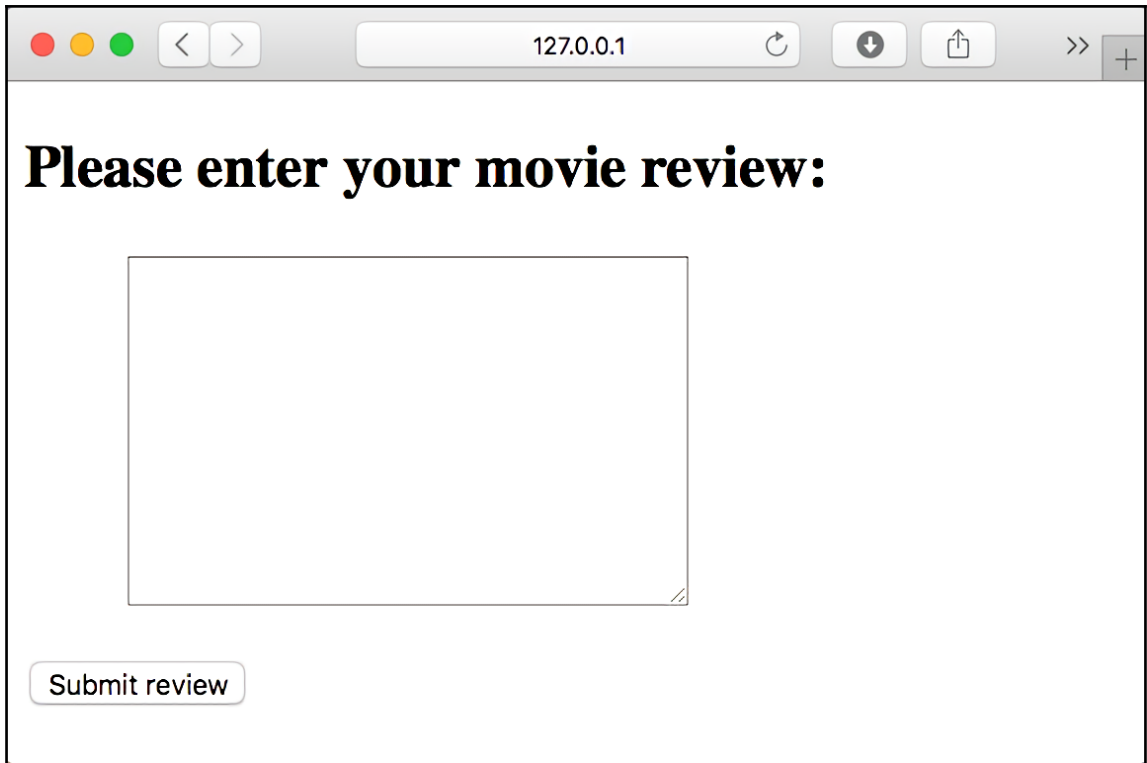
Chapter 8: Applying Machine Learning to Sentiment Analysis

	review	sentiment
0	In 1974, the teenager Martha Moxley (Maggie Gr...	1
1	OK... so... I really like Kris Kristofferson a...	0
2	***SPOILER*** Do not read this, if you think a...	0

Chapter 9: Embedding a Machine Learning Model into a Web Application




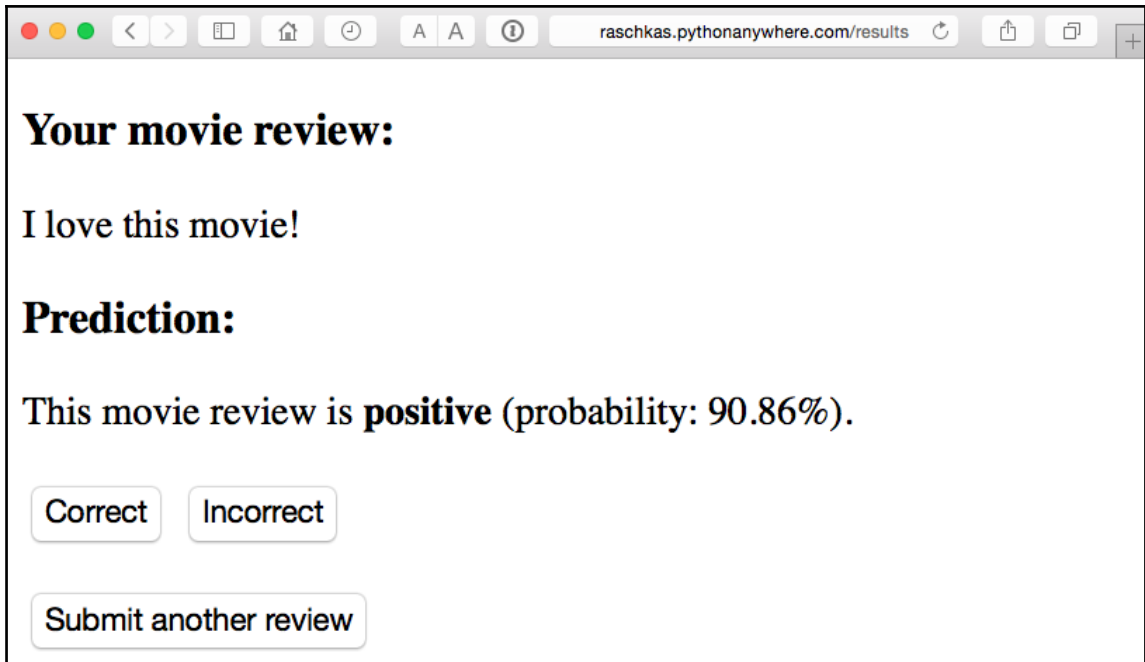


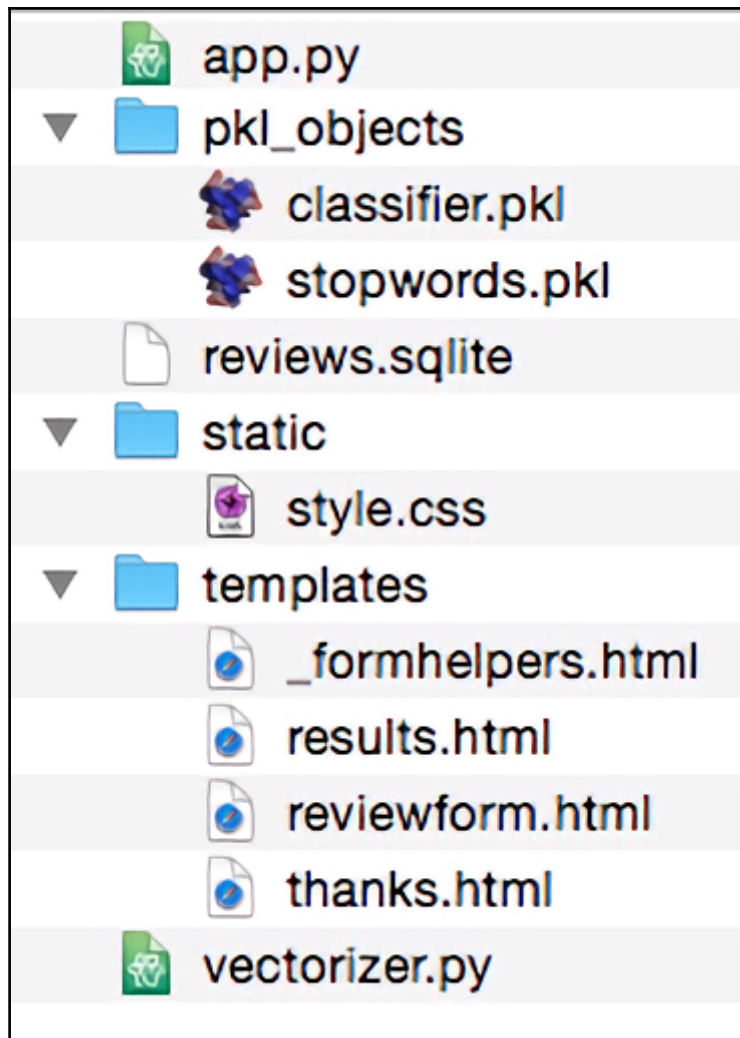


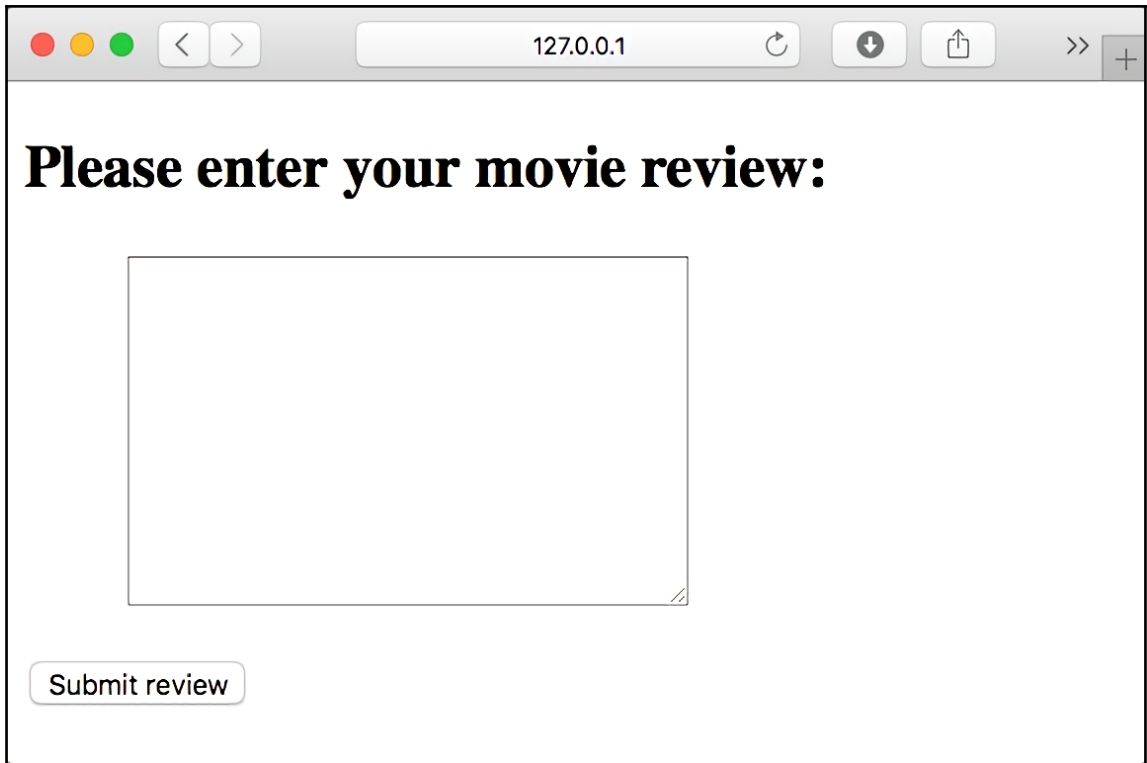
A screenshot of a web browser window. The address bar shows the URL "127.0.0.1". The page content includes the heading "Please enter your movie review:", a large empty text input area, and a "Submit review" button.

Please enter your movie review:

Submit review







A screenshot of a web browser window. The address bar shows the URL "127.0.0.1". The page content includes a heading "Please enter your movie review:", a large empty text input area, and a "Submit review" button.

Please enter your movie review:

Submit review

pythonanywhere.com

pythonanywhere [Send feedback](#) [Forums](#) [Help](#) [Blog](#) [Dashboard](#) [Account](#) [Log out](#)

Consoles **Files** **Web** **Schedule** **Databases**

/ home / **raschkas** / **movieclassifier** [Open Bash console here](#) **3% full (16.4 MB of your 512.0 MB quota)**

Directories

Enter new directory name New directory

- [__pycache__/](#)
- [pkl_objects/](#)
- [static/](#)
- [templates/](#)

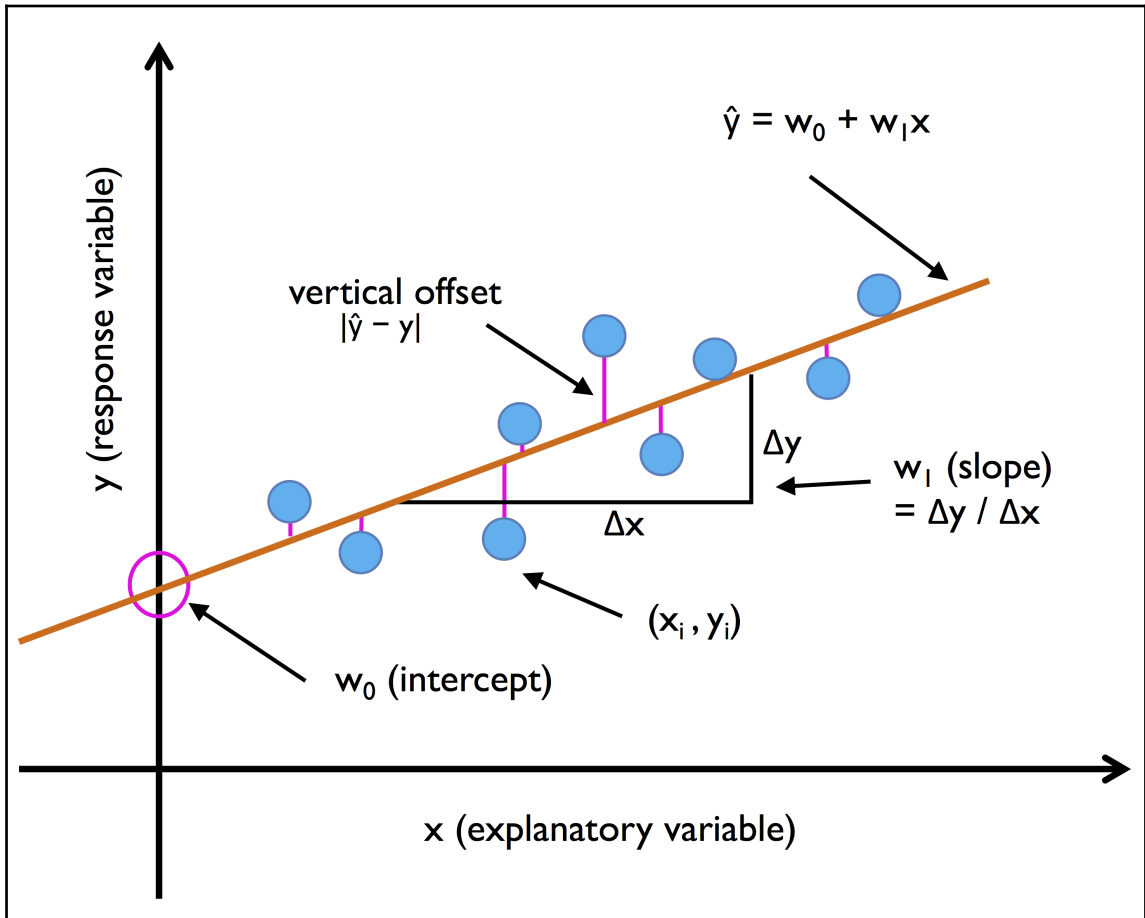
Files


Enter new file name, eg hello.py New file

app.py		2015-07-15 02:21	2.8 KB
reviews.sqlite		2017-04-24 07:57	219.0 KB
vectorizer.py		2015-07-15 00:51	764 bytes

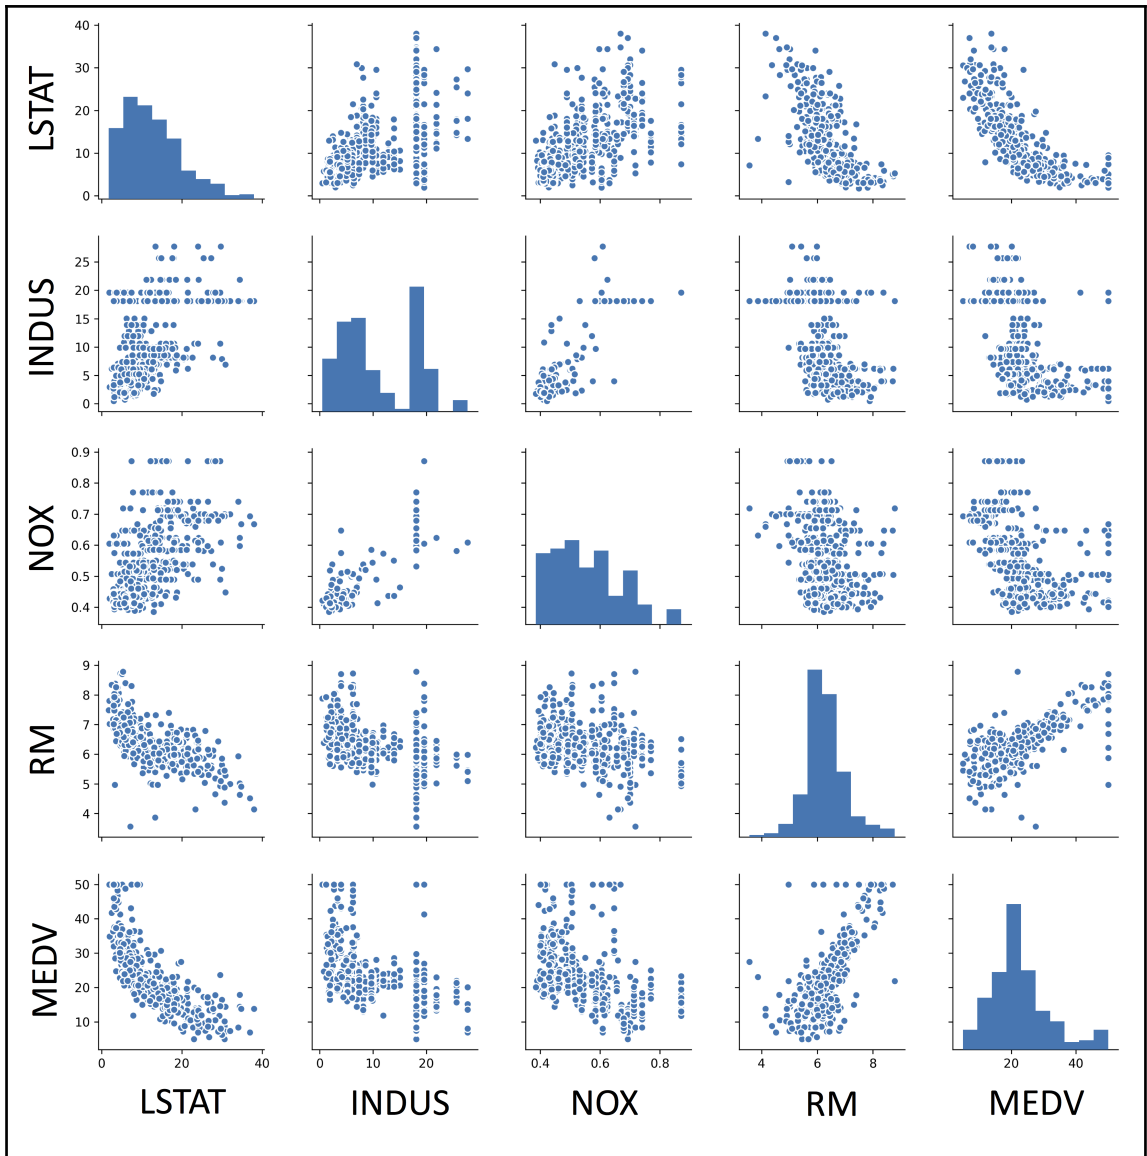
[Upload a file](#)


Chapter 10: Predicting Continuous Target Variables with Regression Analysis








	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	MEDV
0	0.00632	18.0	2.31	0	0.538	6.575	65.2	4.0900	1	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0	0.469	6.421	78.9	4.9671	2	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0	0.469	7.185	61.1	4.9671	2	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0	0.458	6.998	45.8	6.0622	3	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0	0.458	7.147	54.2	6.0622	3	222.0	18.7	396.90	5.33	36.2

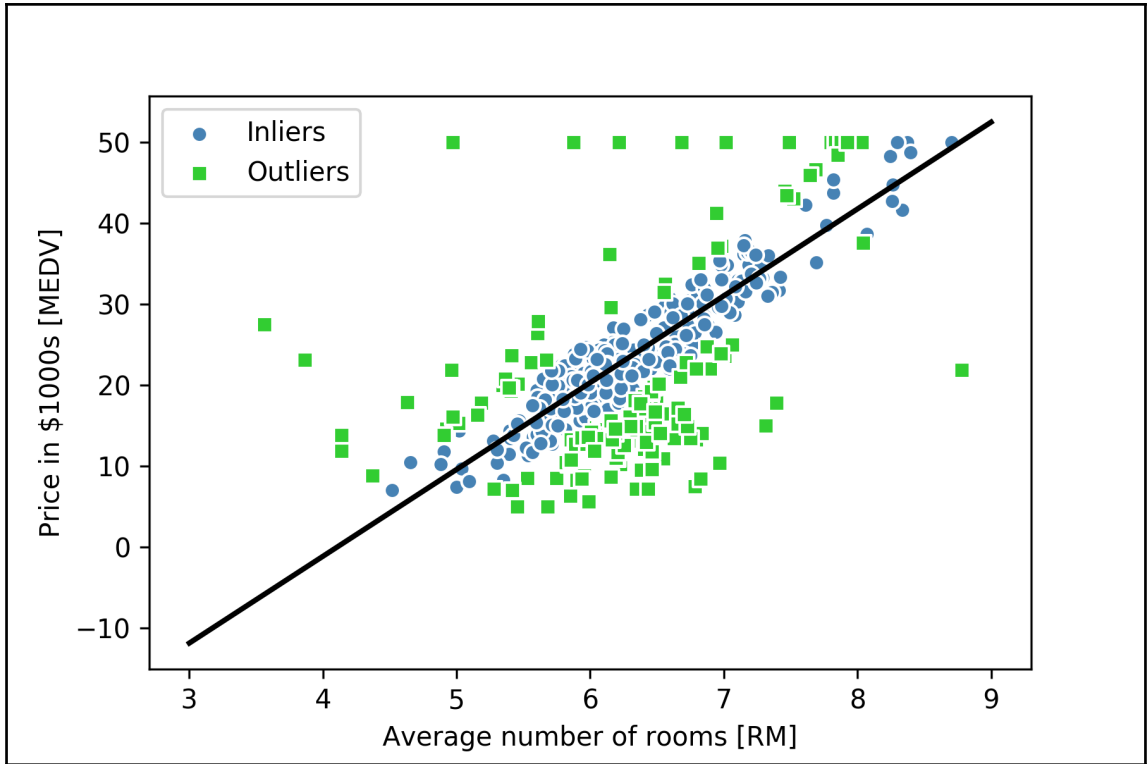





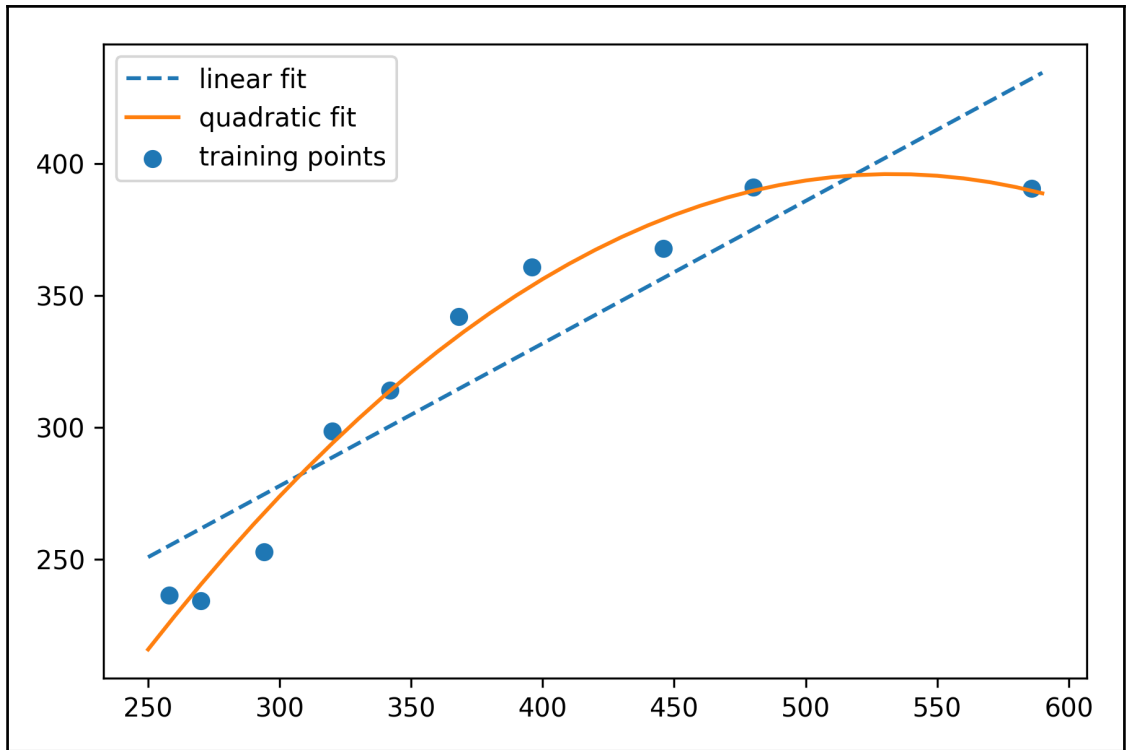








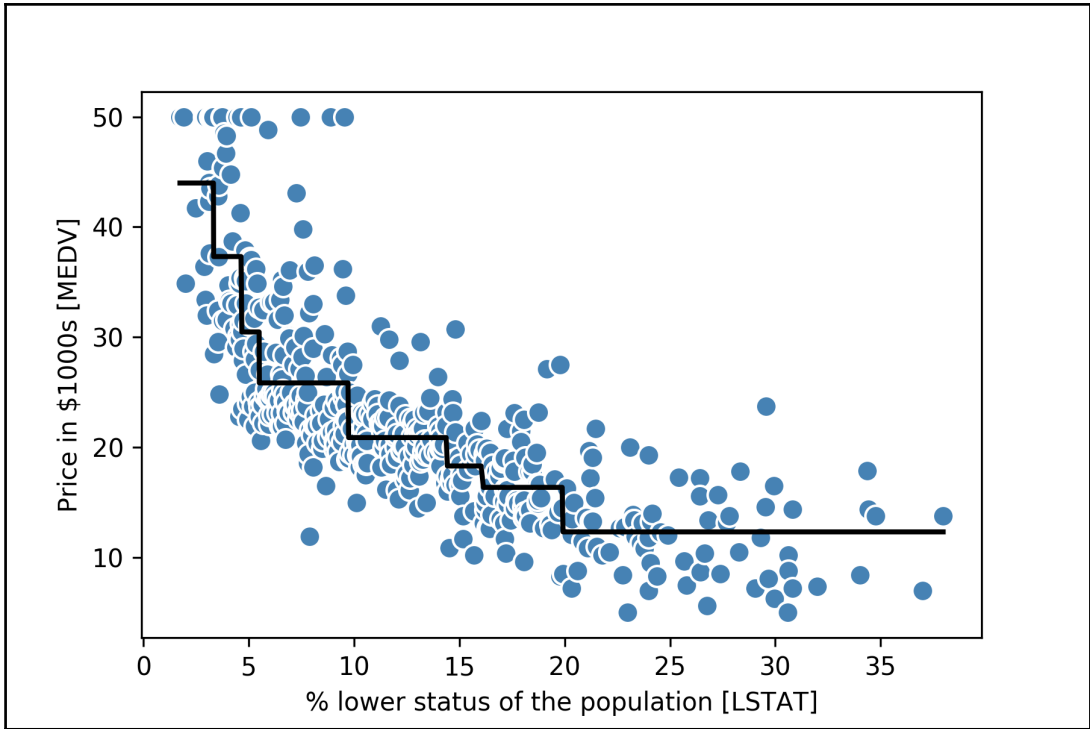




















Chapter 11: Working with Unlabeled Data – Clustering Analysis













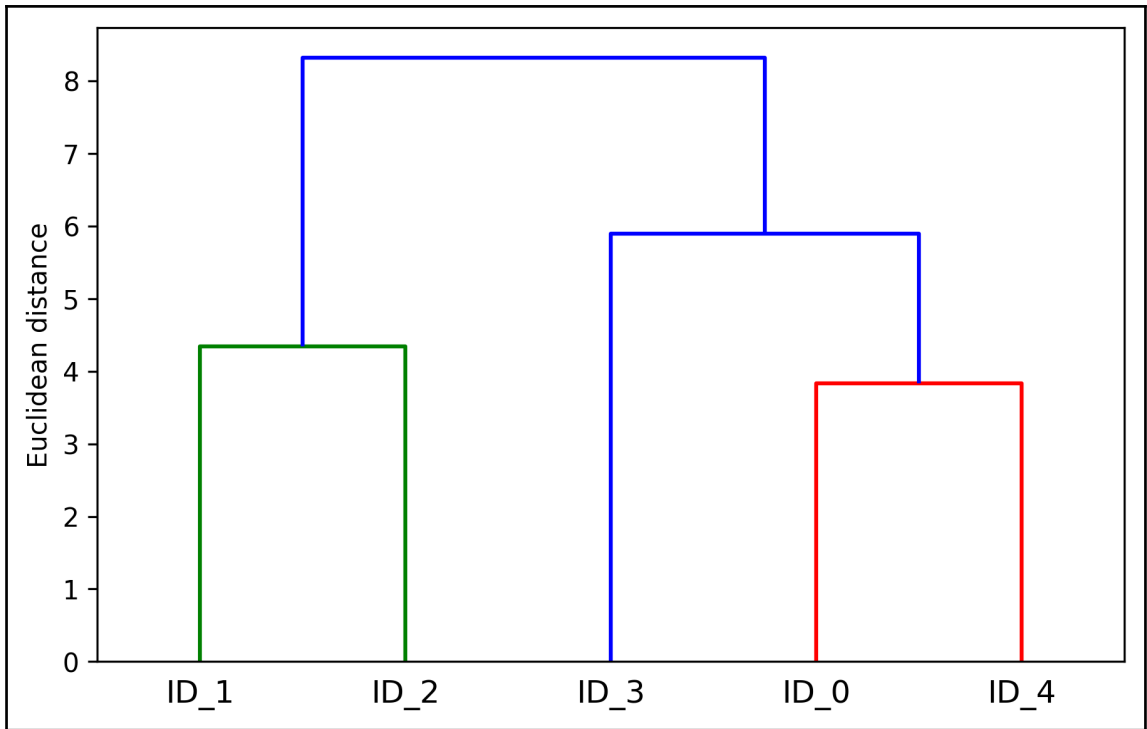





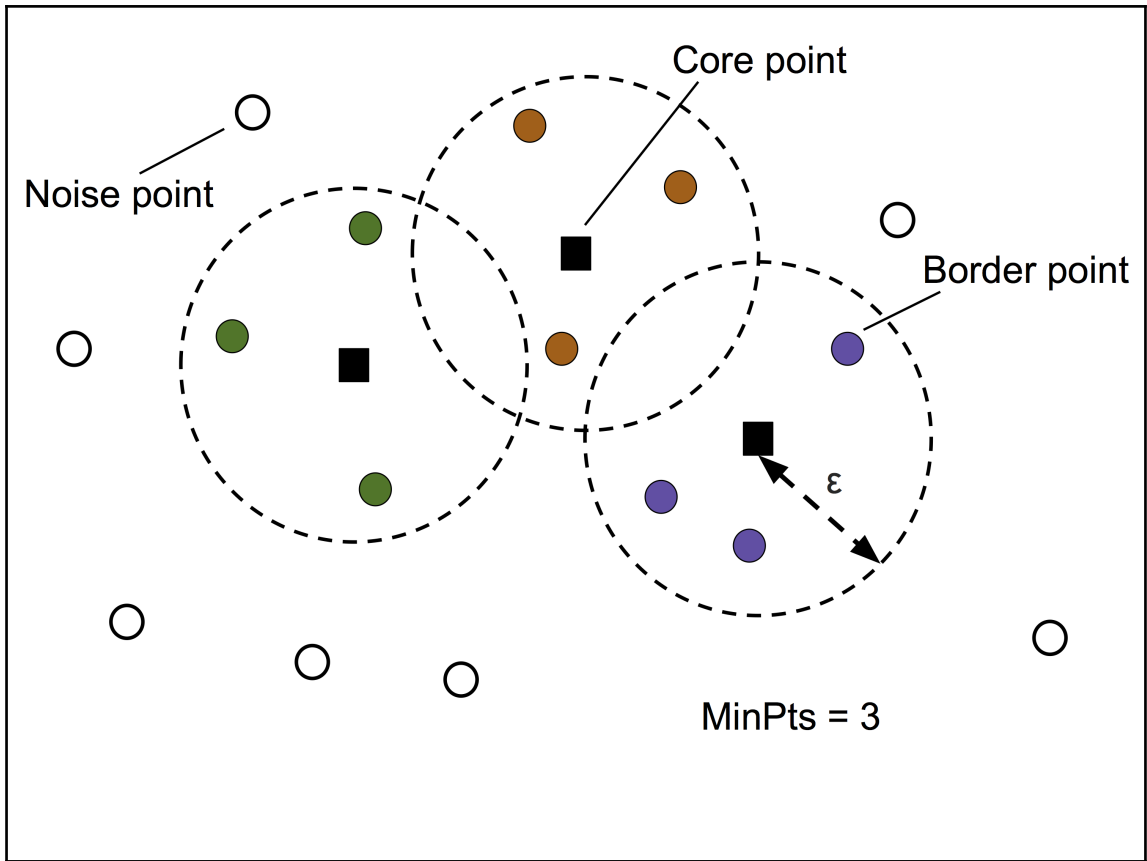
	X	Y	Z
ID_0	6.964692	2.861393	2.268515
ID_1	5.513148	7.194690	4.231065
ID_2	9.807642	6.848297	4.809319
ID_3	3.921175	3.431780	7.290497
ID_4	4.385722	0.596779	3.980443



	X	Y	Z
ID_0	6.964692	2.861393	2.268515
ID_1	5.513148	7.194690	4.231065
ID_2	9.807642	6.848297	4.809319
ID_3	3.921175	3.431780	7.290497
ID_4	4.385722	0.596779	3.980443

	row label 1	row label 2	distance	no. of items in clust.
cluster 1	0.0	4.0	3.835396	2.0
cluster 2	1.0	2.0	4.347073	2.0
cluster 3	3.0	5.0	5.899885	3.0
cluster 4	6.0	7.0	8.316594	5.0






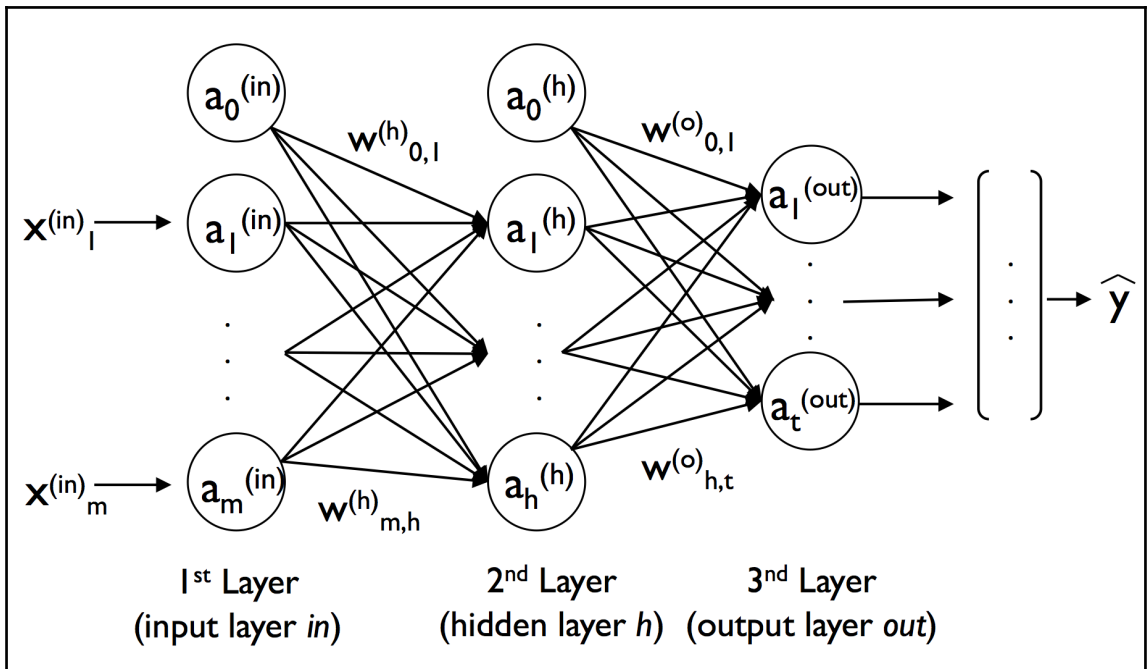
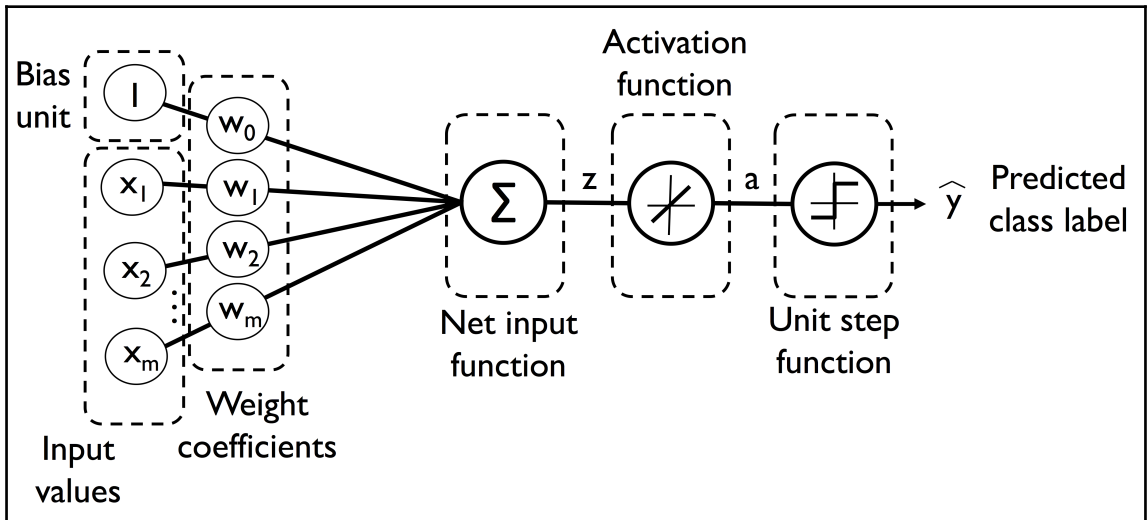


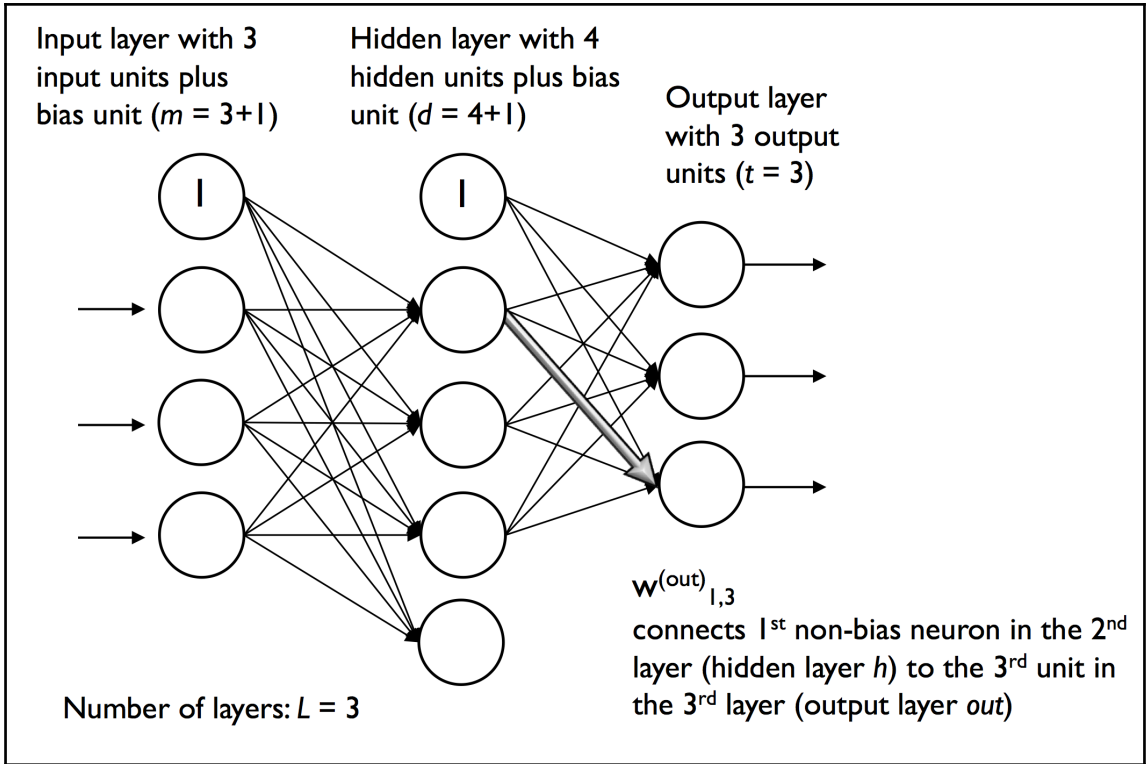



>





Chapter 12: Implementing a Multilayer Artificial Neural Network from Scratch

















1) t: 5 p: 6

5

2) t: 4 p: 9

4

3) t: 4 p: 2

4

4) t: 6 p: 0

6

5) t: 2 p: 7

2

6) t: 5 p: 3

5

7) t: 3 p: 7

3

8) t: 6 p: 0

6

9) t: 3 p: 5

3

10) t: 8 p: 0

8

11) t: 7 p: 1

7

12) t: 3 p: 7

3

13) t: 1 p: 8

1

14) t: 2 p: 6

2

15) t: 2 p: 8

2

16) t: 7 p: 3

7

17) t: 8 p: 4

8

18) t: 5 p: 8

5

19) t: 4 p: 9

4

20) t: 9 p: 7

9

21) t: 2 p: 7

2

22) t: 3 p: 5

3

23) t: 8 p: 9

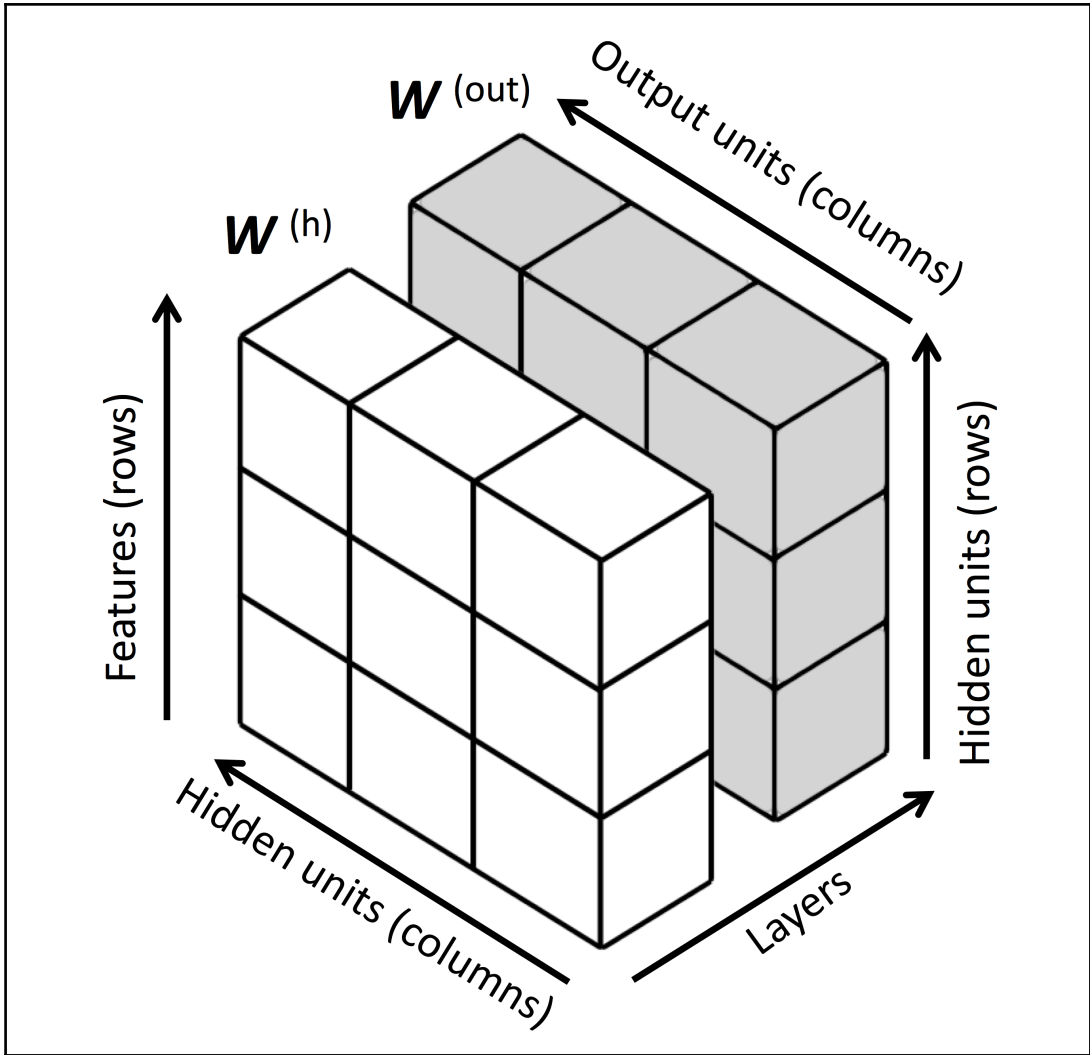
8


24) t: 5 p: 4

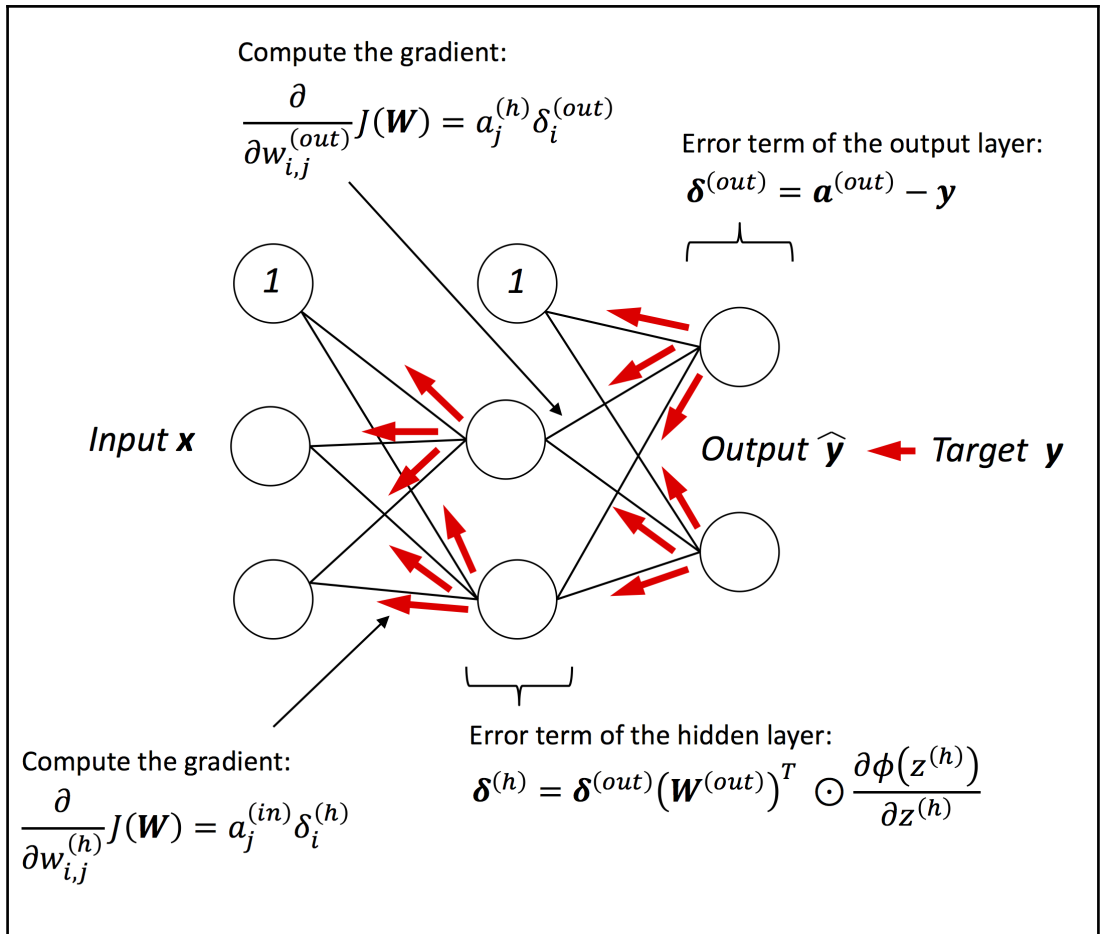
5

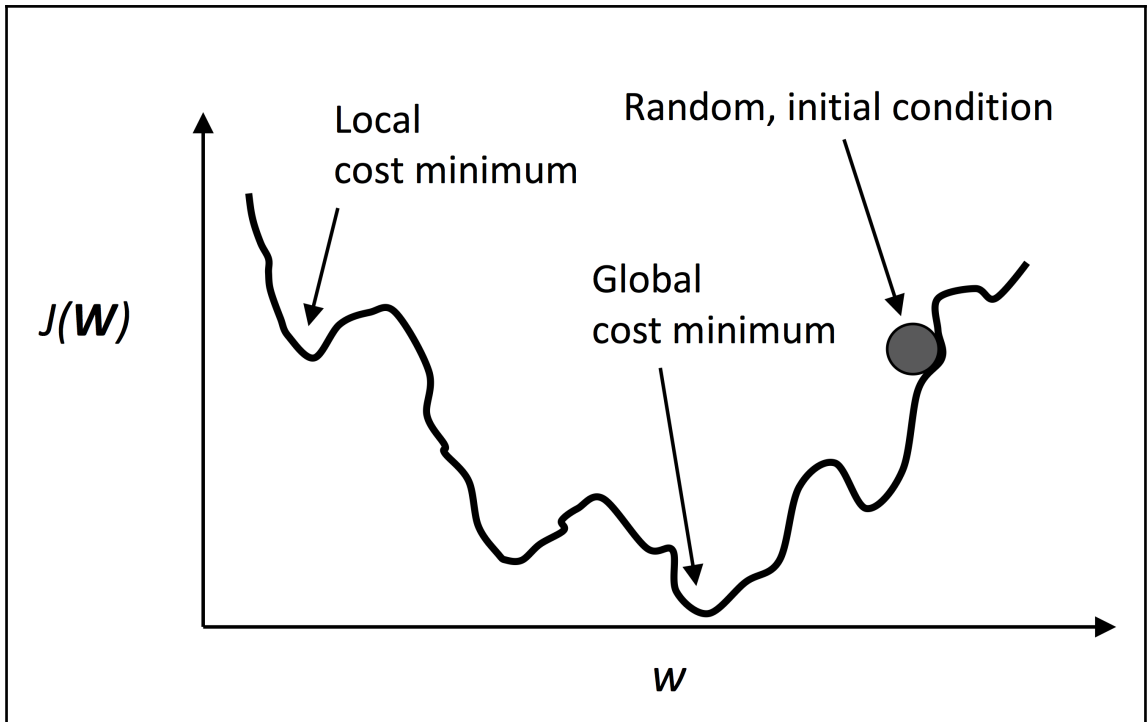
25) t: 1 p: 2

1










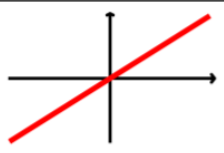
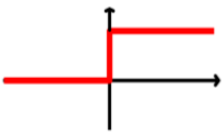
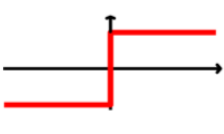
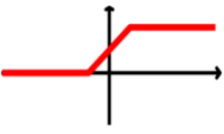
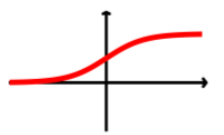
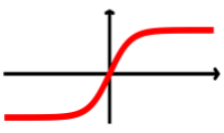
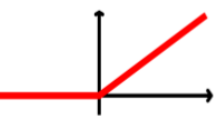


Chapter 13: Parallelizing Neural Network Training with TensorFlow

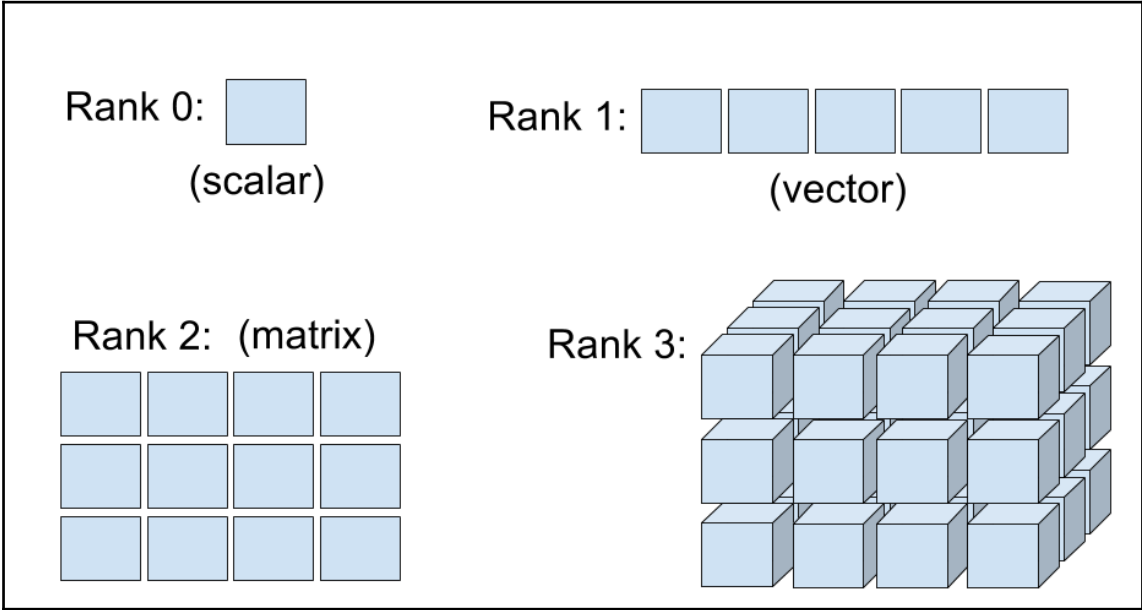
Specifications	Intel® Core™ i7-6900K Processor Extreme Ed.	NVIDIA GeForce® GTX™ 1080 Ti
Base Clock Frequency	3.2 GHz	< 1.5 GHz
Cores	8	3584
Memory Bandwidth	64 GB/s	484 GB/s
Floating-Point Calculations	409 GFLOPS	11300 GFLOPS
Cost	~ \$1000.00	~ \$700.00



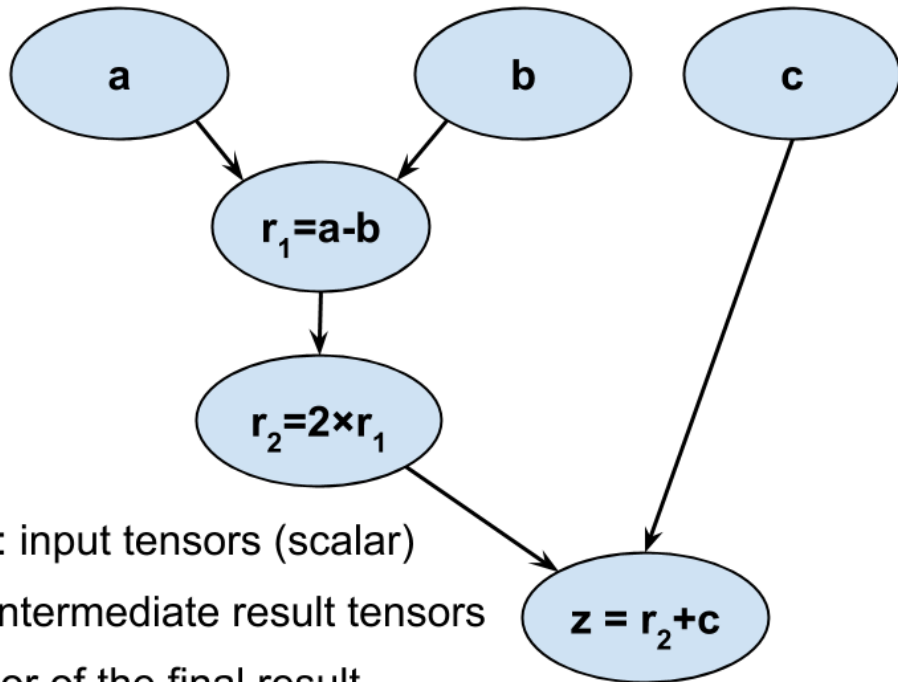


Activation Function	Equation	Example	1D Graph
Linear	$\phi(z) = z$	Adaline, linear regression	
Unit Step (Heaviside Function)	$\phi(z) = \begin{cases} 0 & z < 0 \\ 0.5 & z = 0 \\ 1 & z > 0 \end{cases}$	Perceptron variant	
Sign (signum)	$\phi(z) = \begin{cases} -1 & z < 0 \\ 0 & z = 0 \\ 1 & z > 0 \end{cases}$	Perceptron variant	
Piece-wise Linear	$\phi(z) = \begin{cases} 0 & z \leq -\frac{1}{2} \\ z + \frac{1}{2} & -\frac{1}{2} \leq z \leq \frac{1}{2} \\ 1 & z \geq \frac{1}{2} \end{cases}$	Support vector machine	
Logistic (sigmoid)	$\phi(z) = \frac{1}{1 + e^{-z}}$	Logistic regression, Multilayer NN	
Hyperbolic Tangent (tanh)	$\phi(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$	Multilayer NN, RNNs	
ReLU	$\phi(z) = \begin{cases} 0 & z < 0 \\ z & z > 0 \end{cases}$	Multilayer NN, CNNs	

Chapter 14: Going Deeper – The Mechanics of TensorFlow




**Computation graph implementing
the equation $z = 2 \times (a - b) + c$**





a, b, c: input tensors (scalar)

r₁, r₂: intermediate result tensors


z: tensor of the final result









Python if



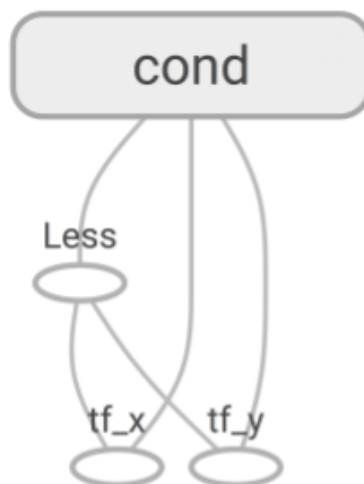
tf.cond(...)

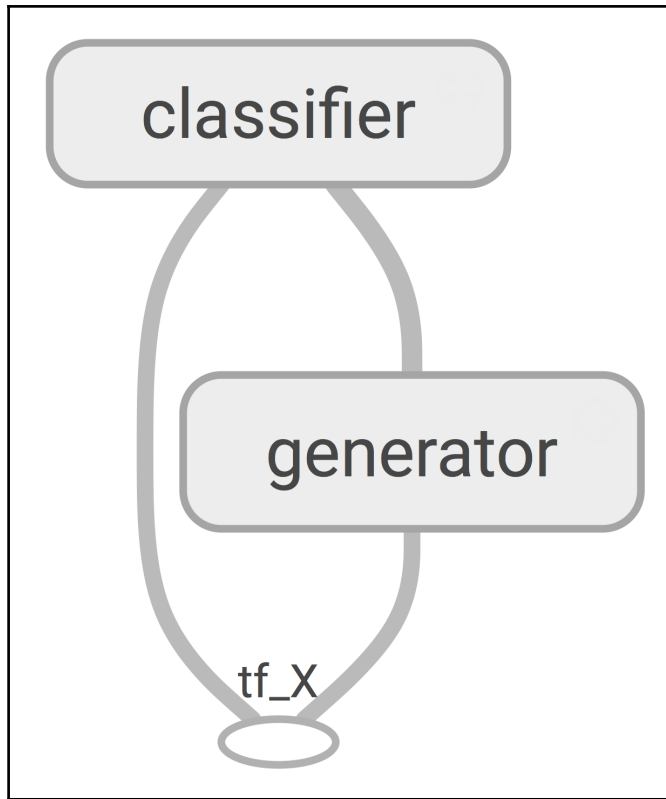


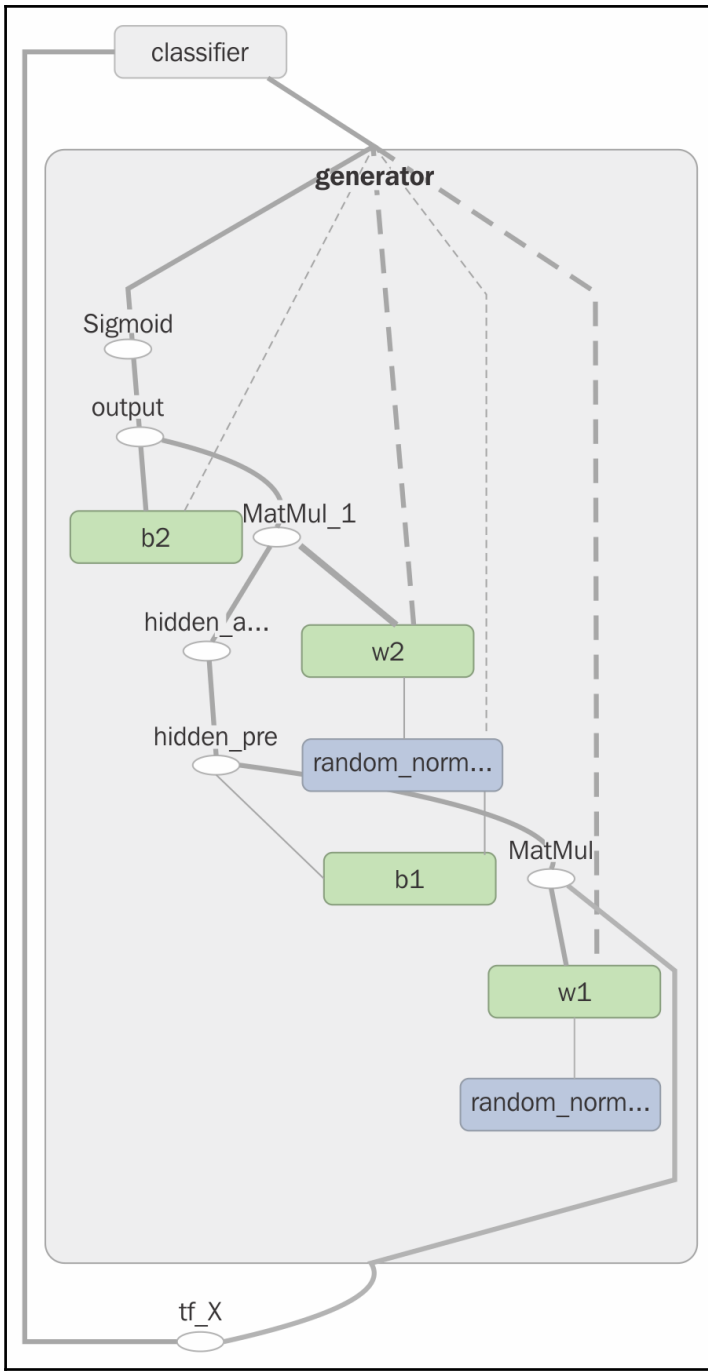
Python if

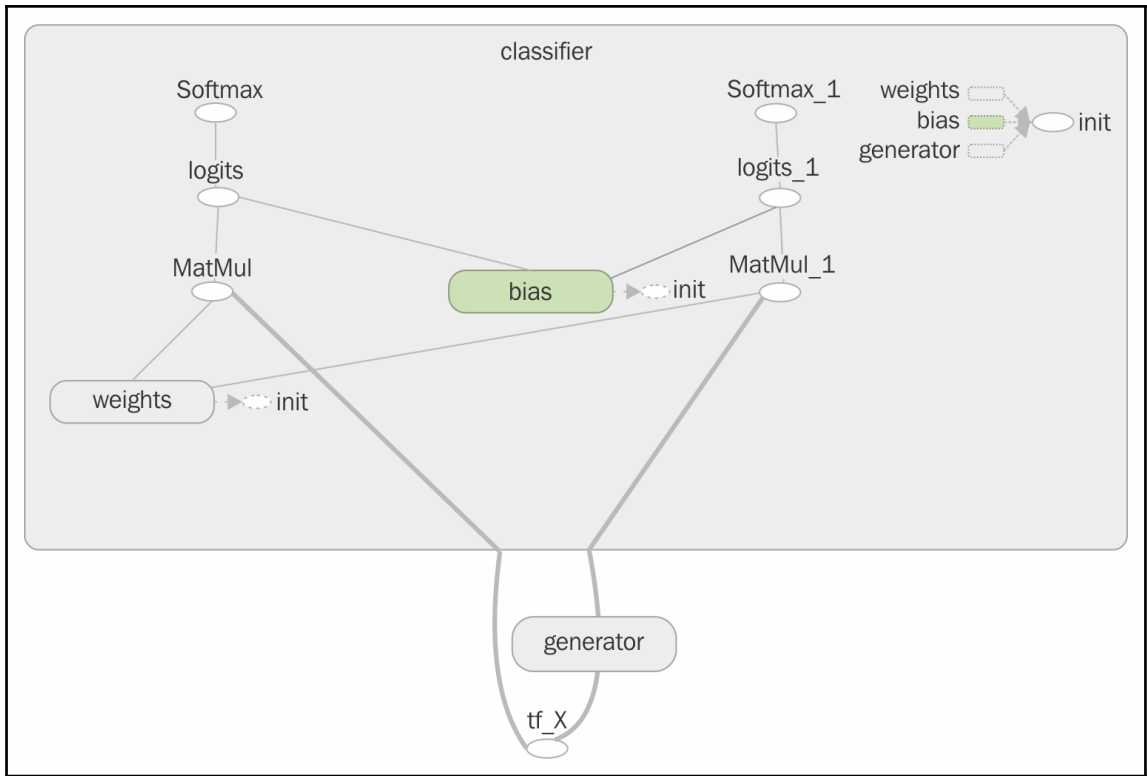


tf.cond(...)

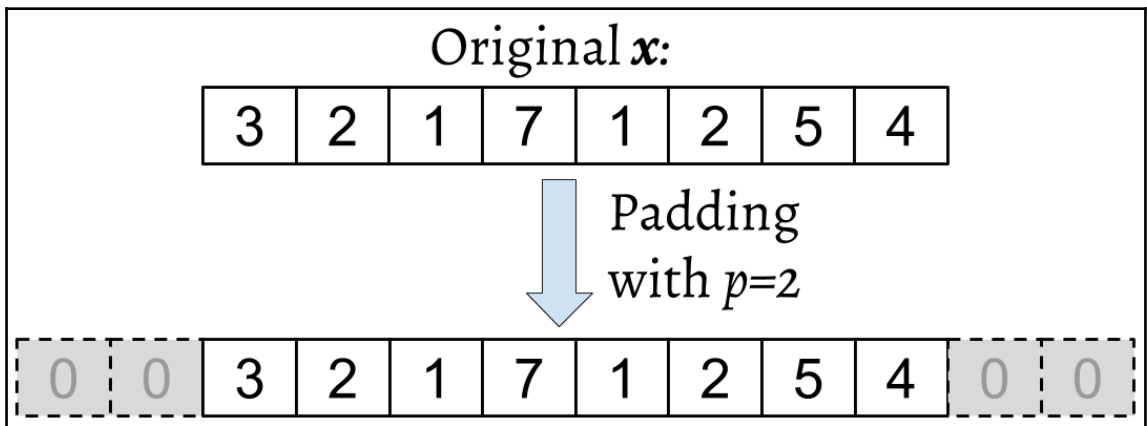



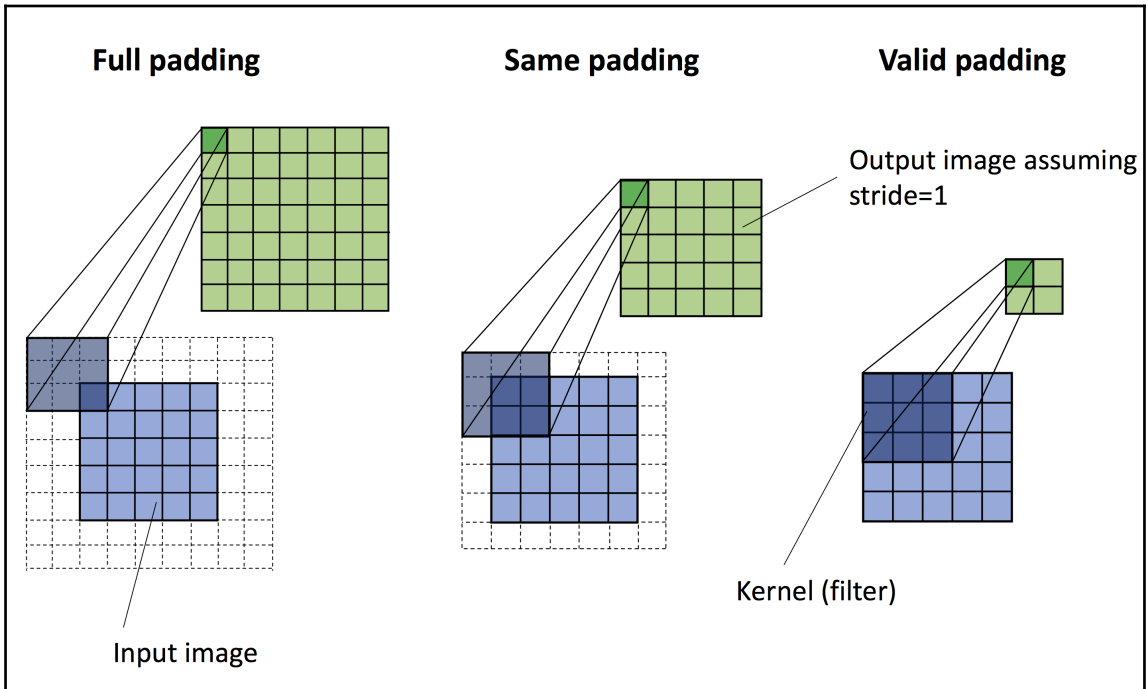
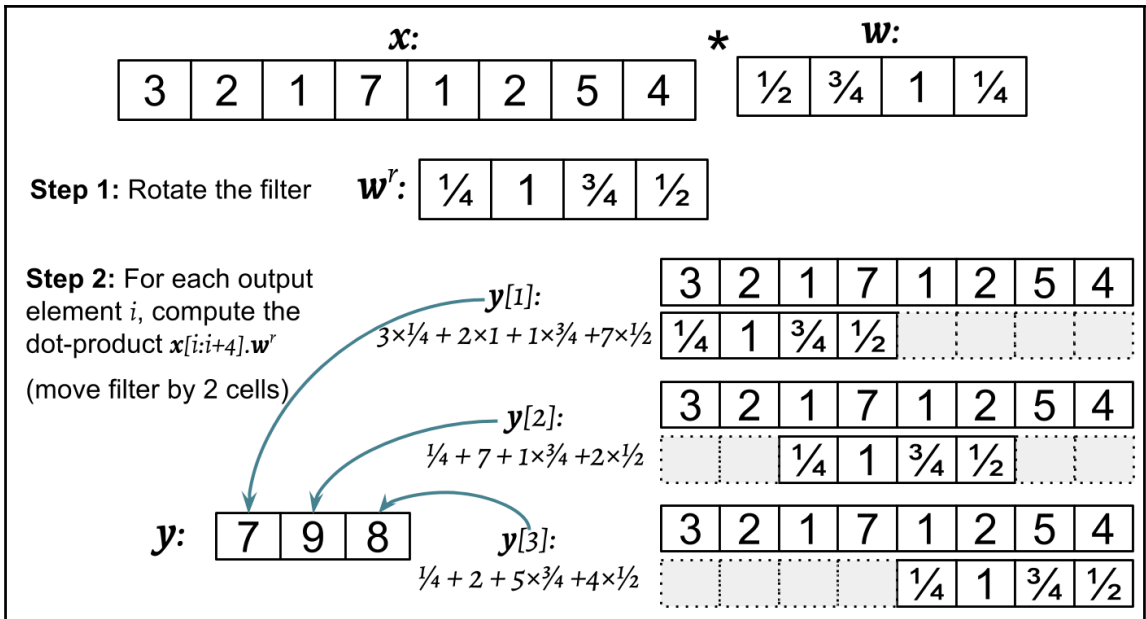



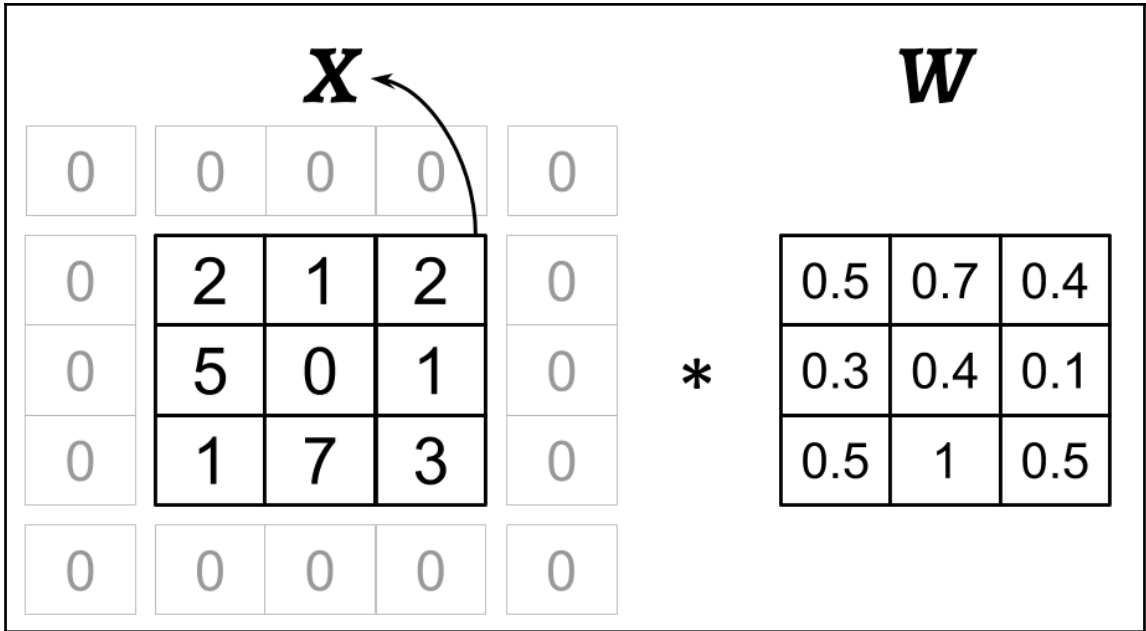


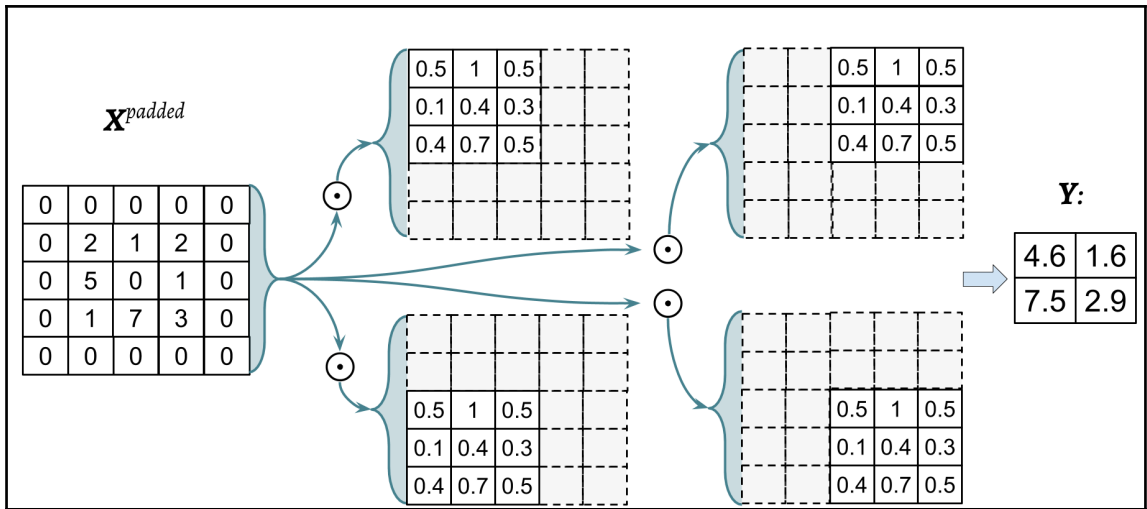


Chapter 15: Classifying Images with Deep Convolutional Neural Networks

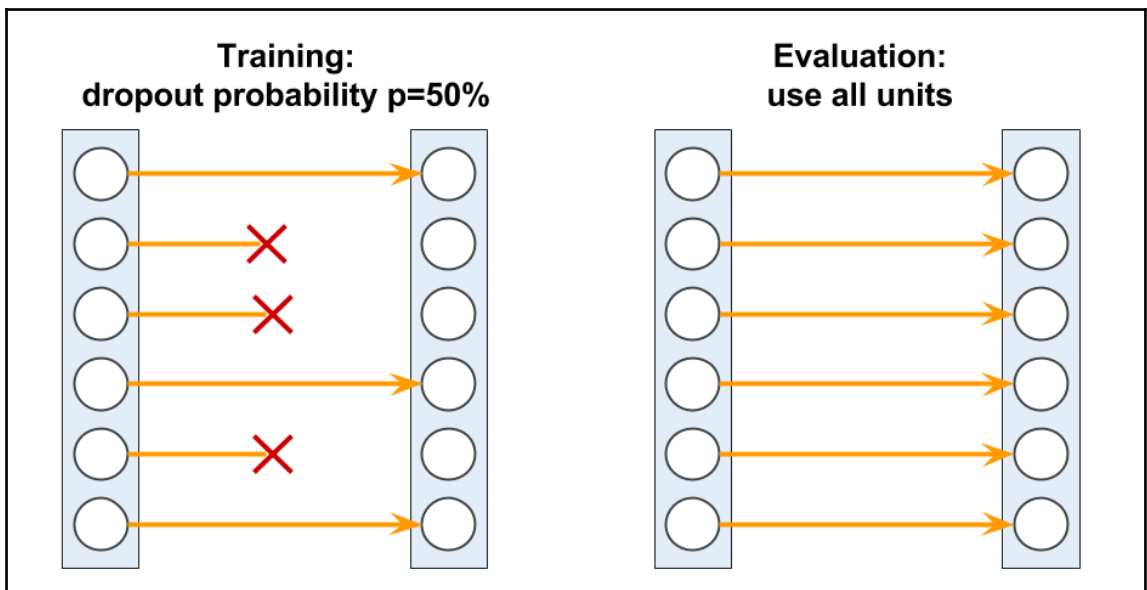








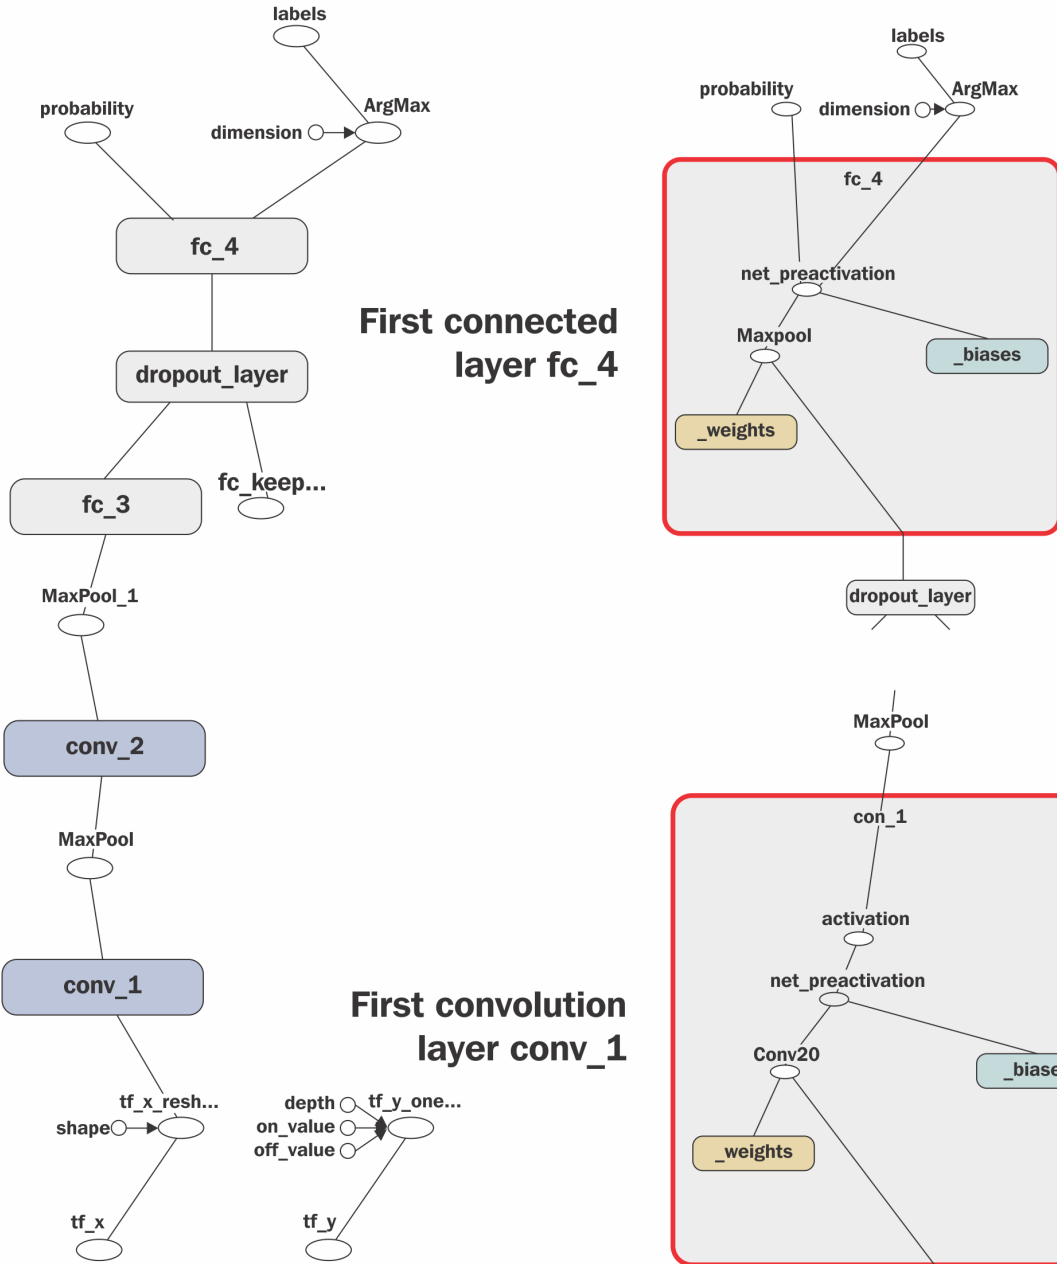


$$\begin{array}{l}
 X_1 = \begin{bmatrix} 10 & 255 & 125 & 0 & 170 & 100 \\ 70 & 255 & 105 & 25 & 25 & 70 \\ 255 & 0 & 150 & 0 & 10 & 10 \\ 0 & 255 & 10 & 10 & 150 & 20 \\ 70 & 15 & 200 & 100 & 95 & 0 \\ 35 & 25 & 100 & 20 & 0 & 60 \end{bmatrix} \\
 X_2 = \begin{bmatrix} 100 & 100 & 100 & 50 & 100 & 50 \\ 95 & 255 & 100 & 125 & 125 & 170 \\ 80 & 40 & 10 & 10 & 125 & 150 \\ 255 & 30 & 150 & 20 & 120 & 125 \\ 30 & 30 & 150 & 100 & 70 & 70 \\ 70 & 30 & 100 & 200 & 70 & 95 \end{bmatrix}
 \end{array}
 \xrightarrow{\text{max-pooling } P_{2 \times 2}}
 \begin{bmatrix} 255 & 125 & 170 \\ 255 & 150 & 150 \\ 70 & 200 & 95 \end{bmatrix}$$









Graph of the multilayer CNN




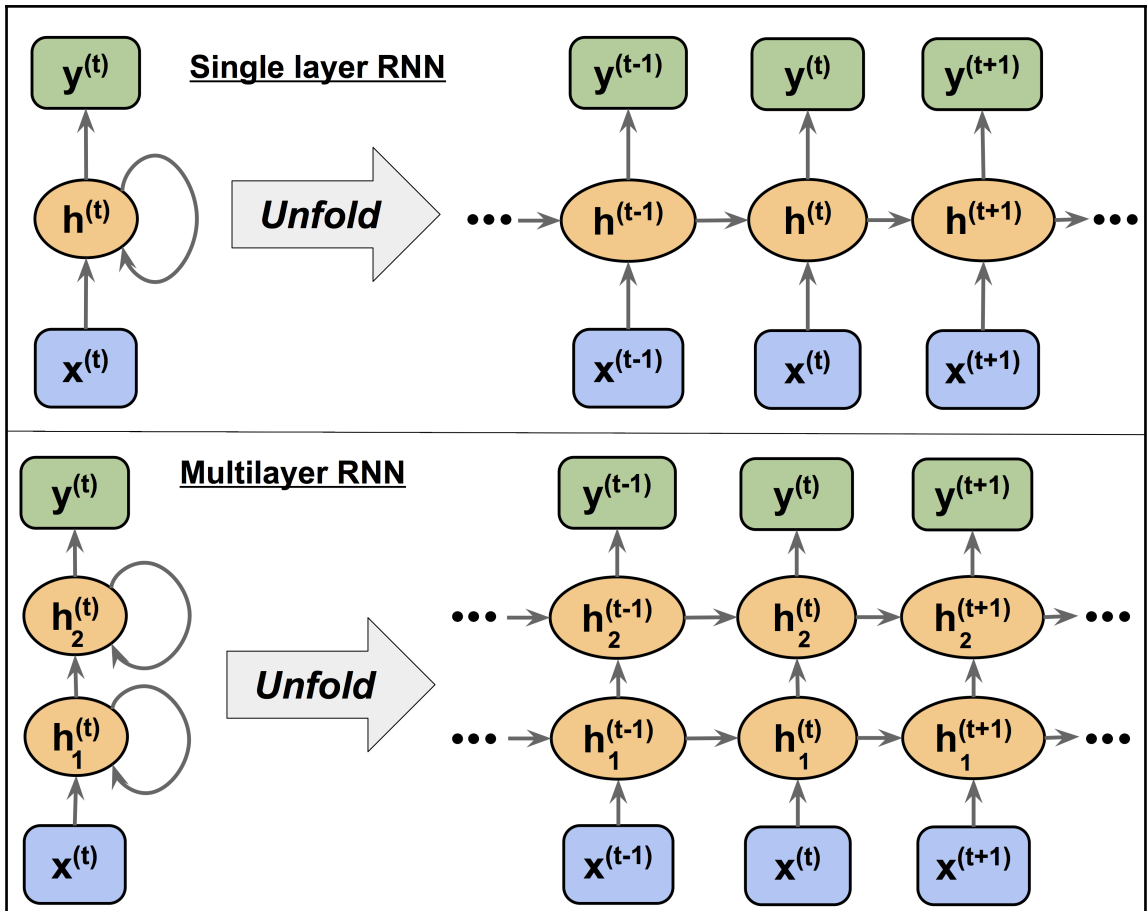
Chapter 16: Modeling Sequential Data Using Recurrent Neural Networks

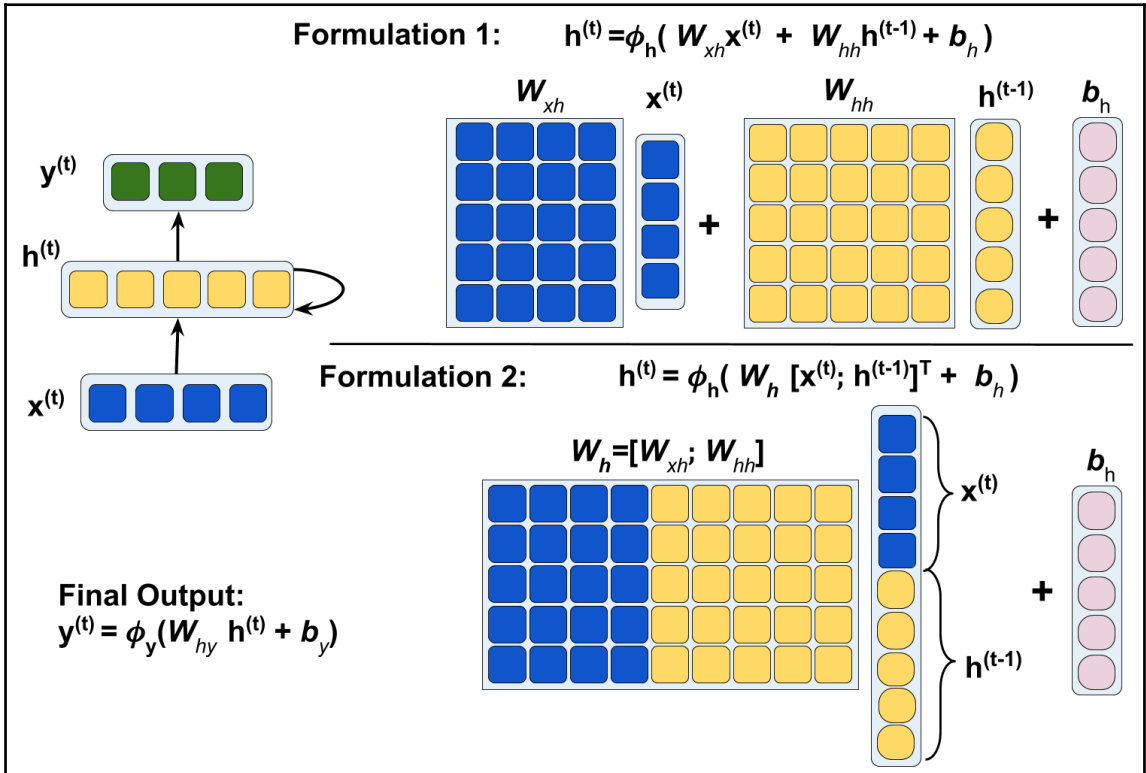
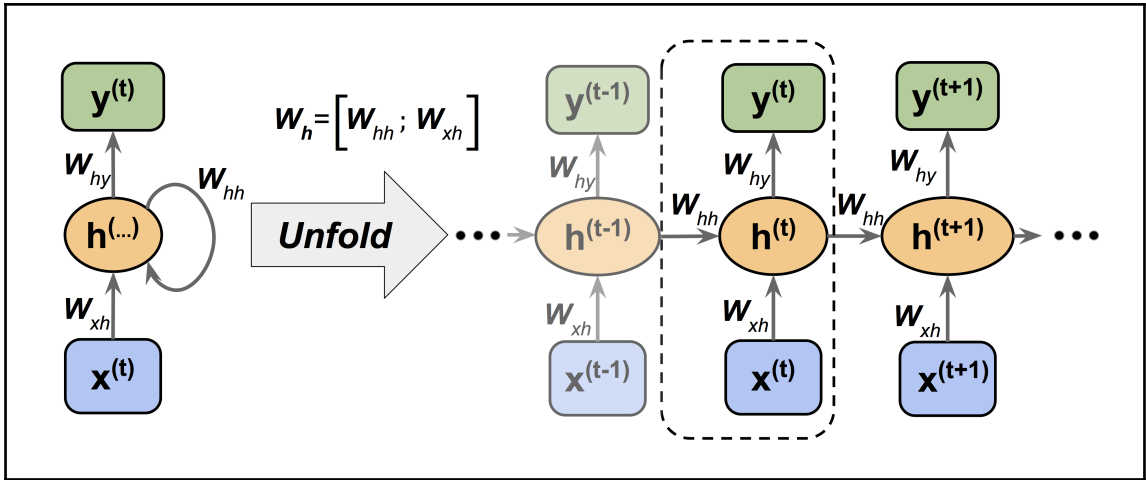


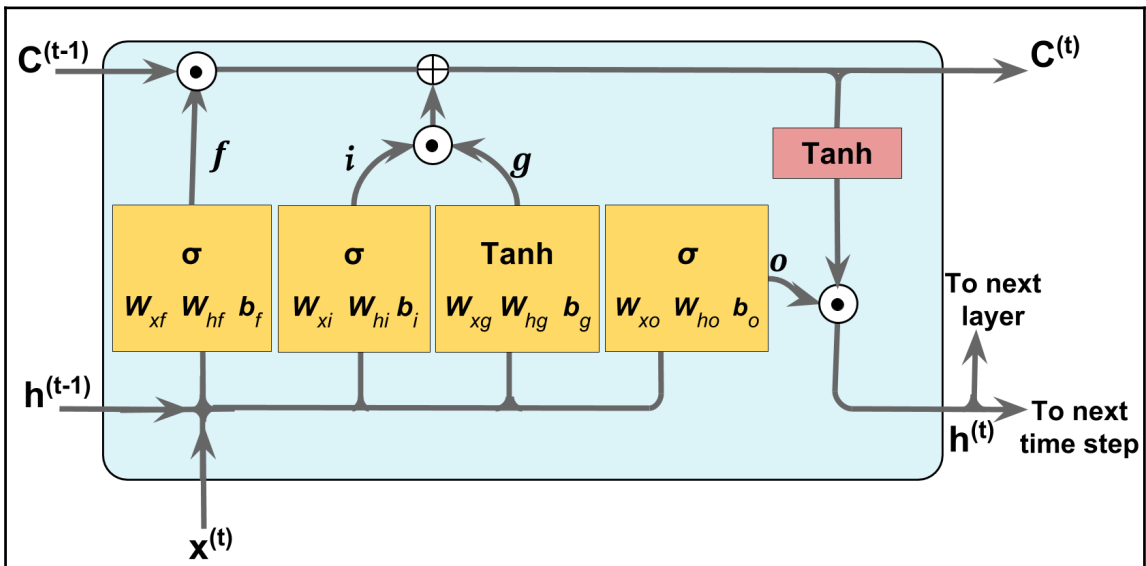
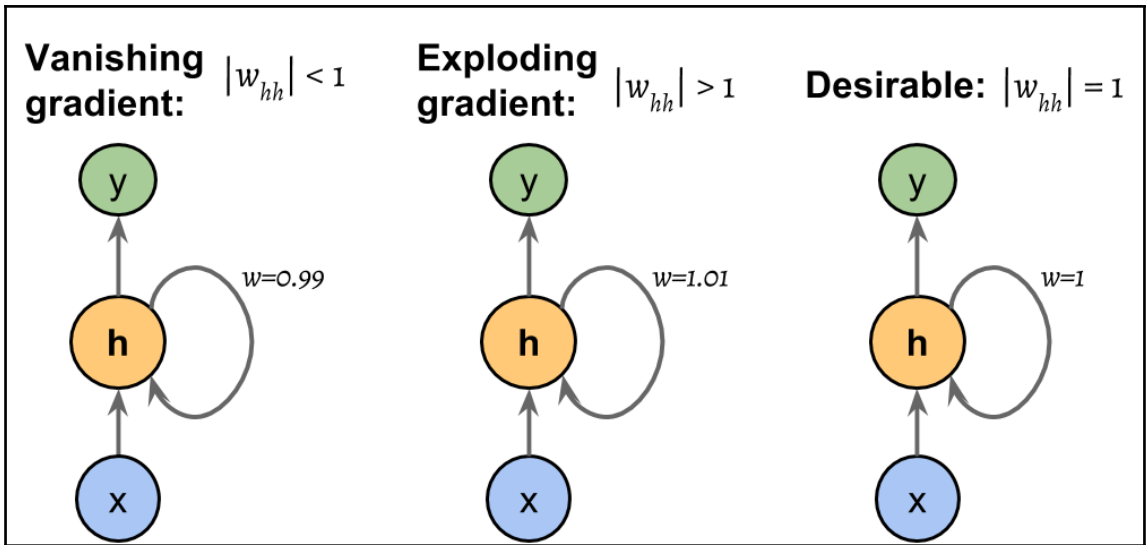
Chapter 16: Modeling Sequential Data Using Recurrent Neural Networks

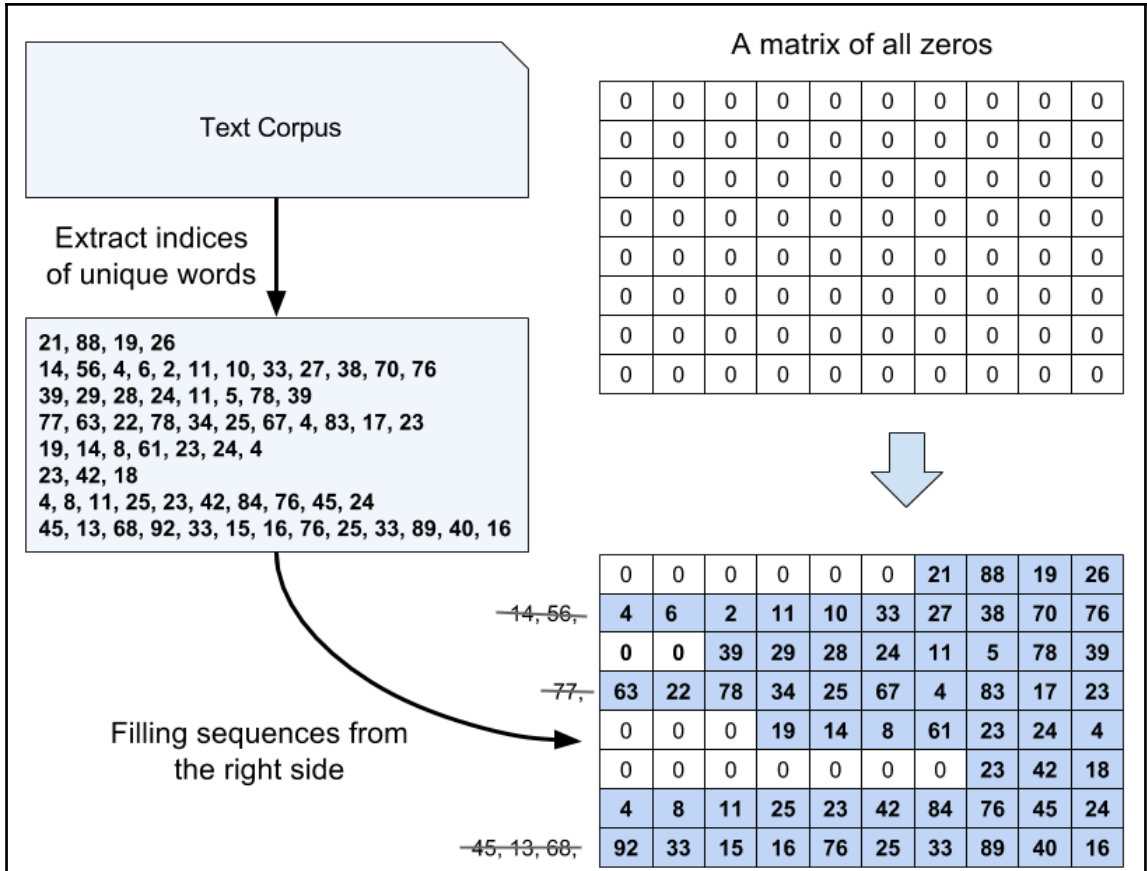


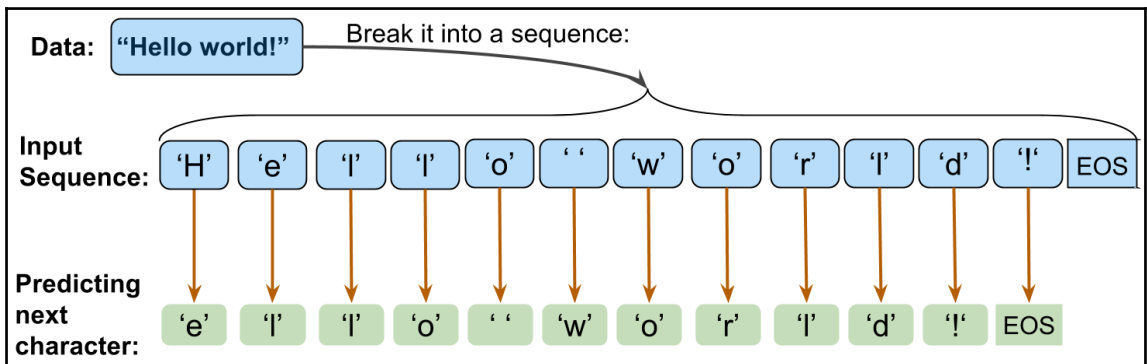
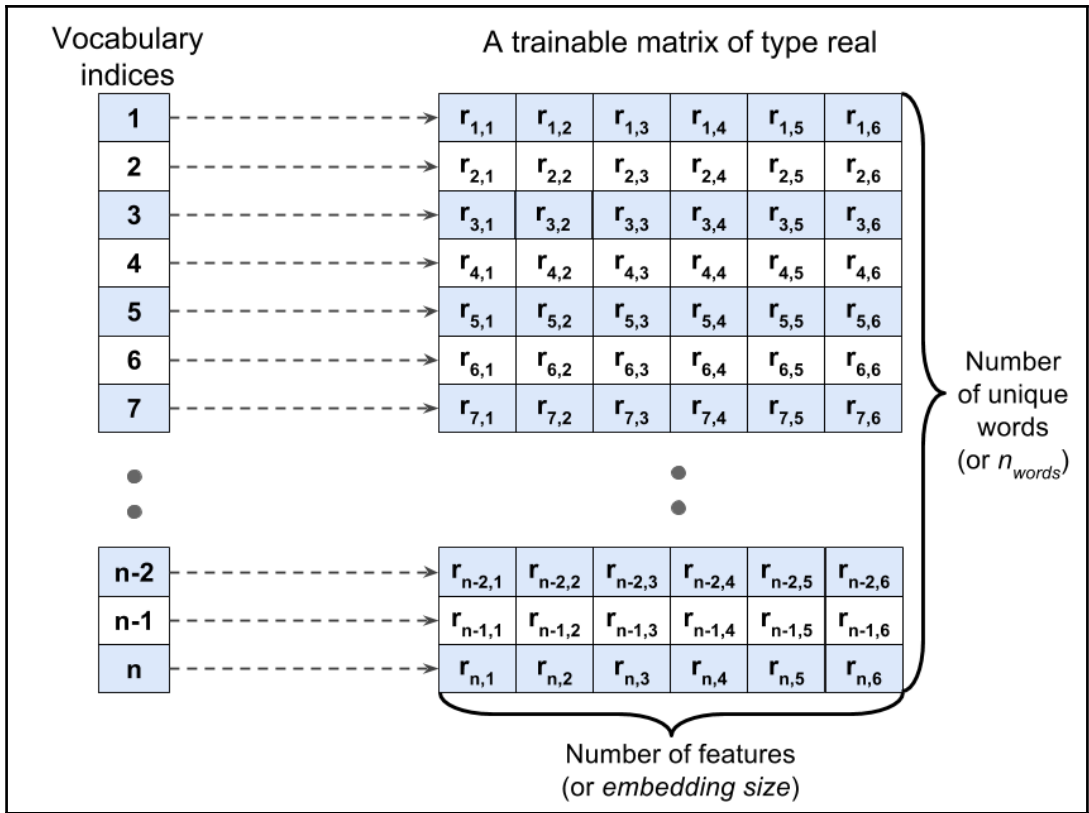


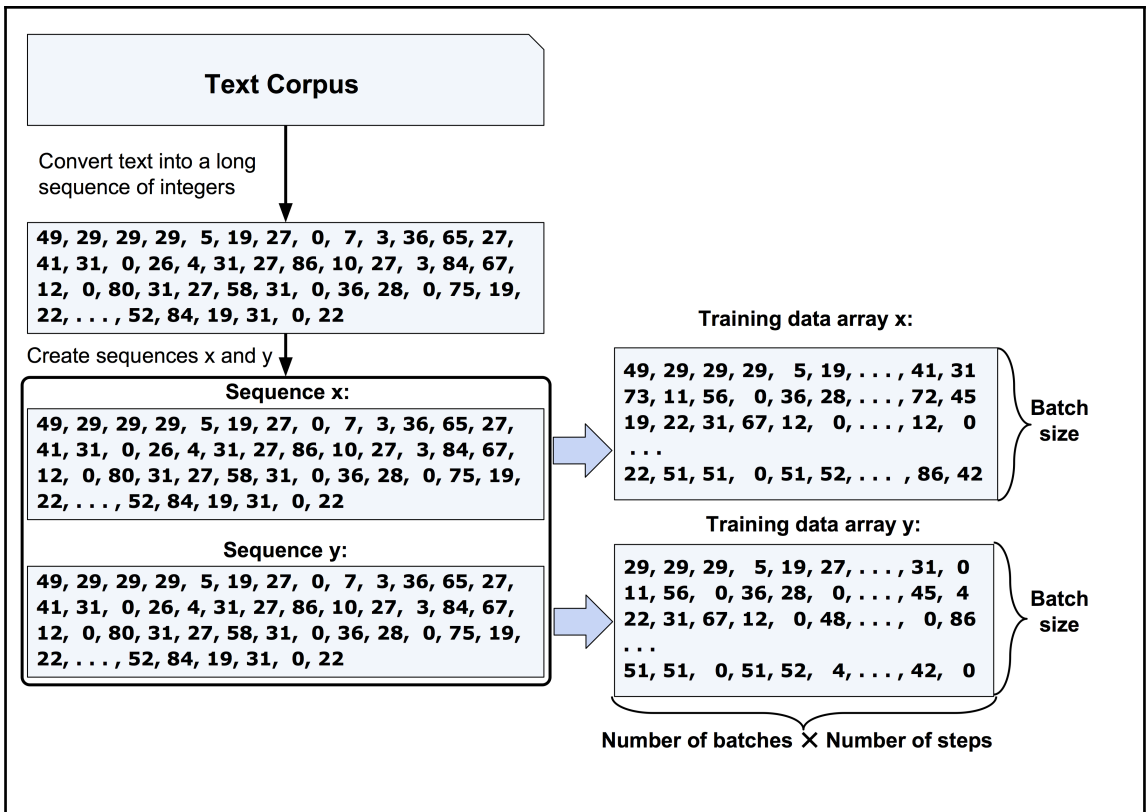
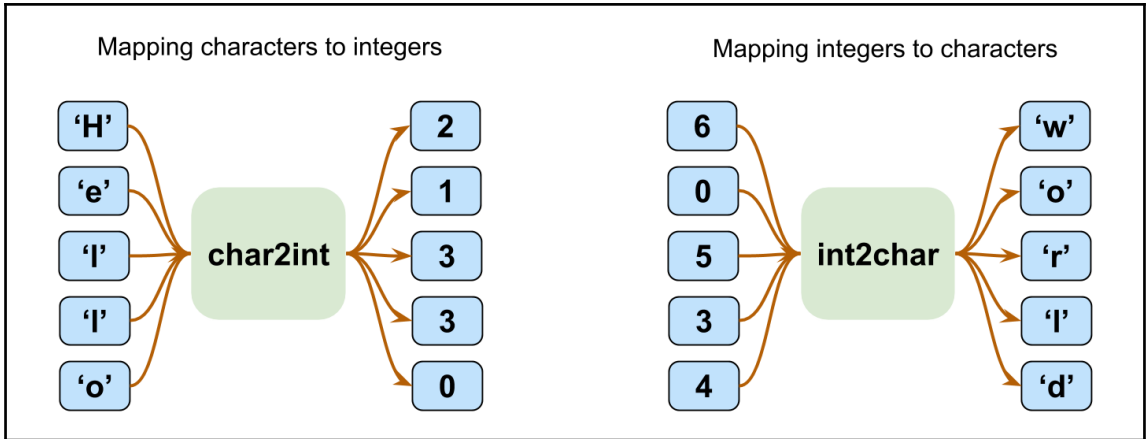


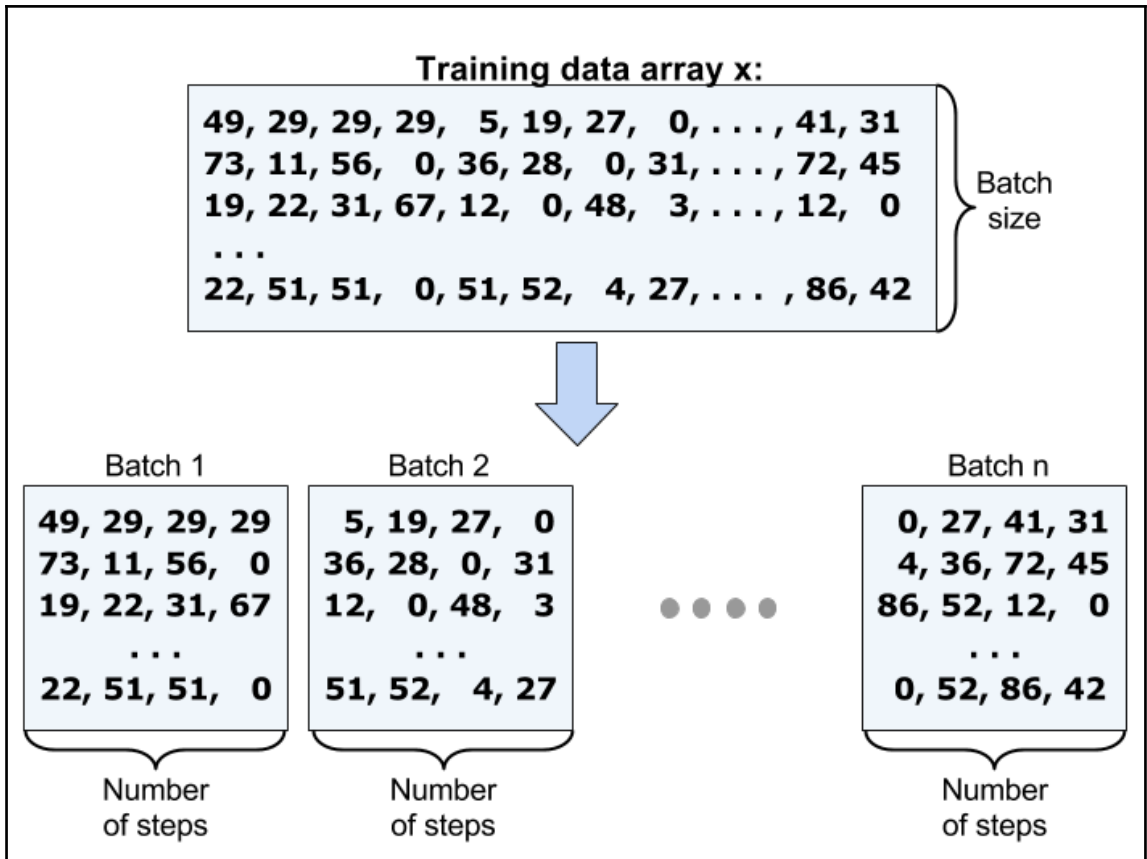












The stall soues tay and the hates,
 The perse in there is that so the meanes this made there

Ham. Ile teath thes are this makere of a driane,
 Why shis mestend the Casst of is singe,
 In this to this, to mers it is for marth,
 Ase hinee sim thig tald ow a tore andere,
 In histhene tistere shere this wile and my Lord:
 And tit mighes the secleer alllost heruen, and that hash to sall and hears,
 If you his moses tonger and mout ofr mesting a forte tis at

Pomin. Where in you dist and sintere shan shall