




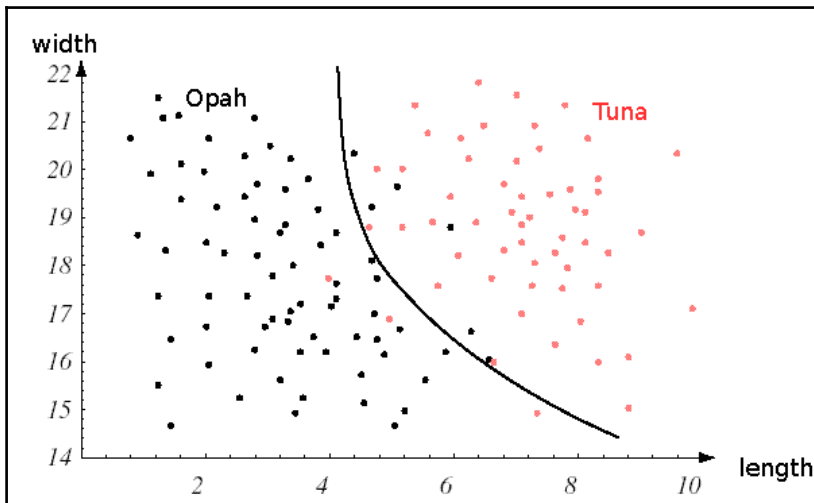
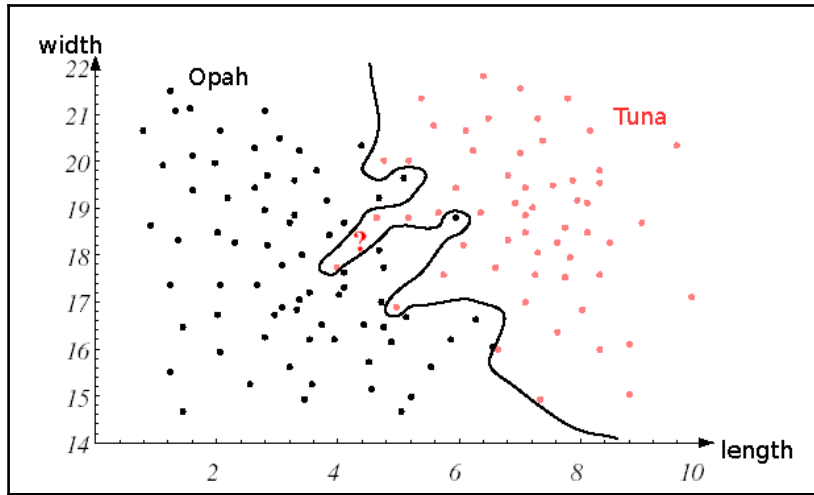
Table of Contents

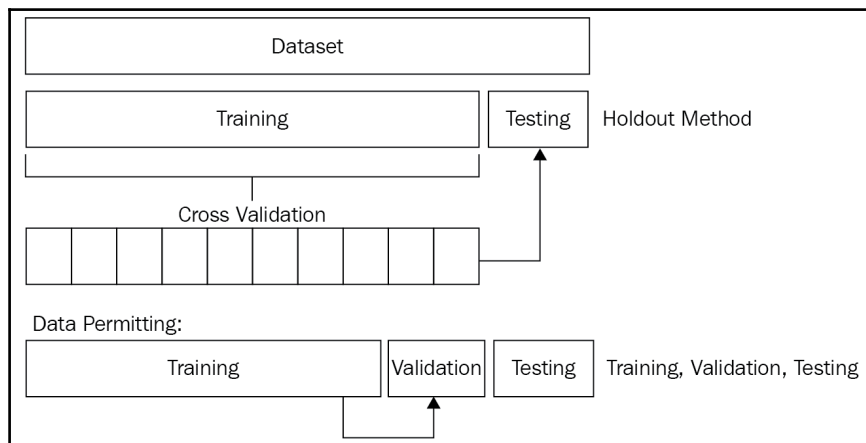
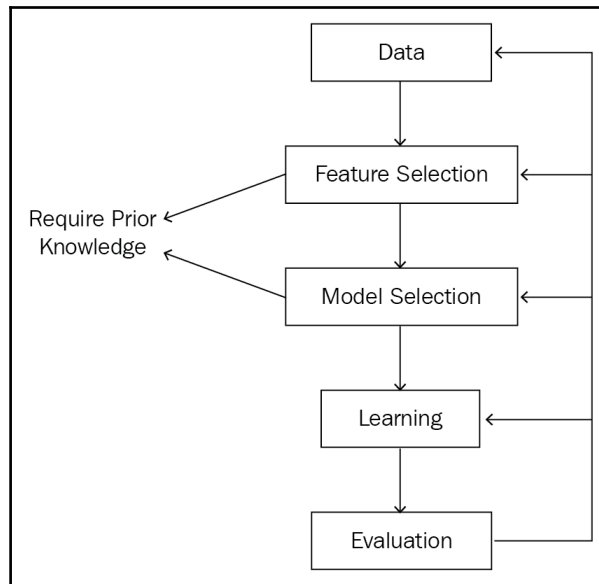
	1
Index	76

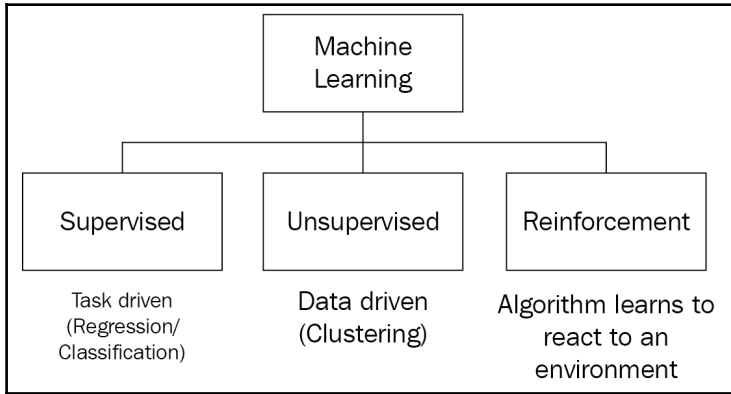
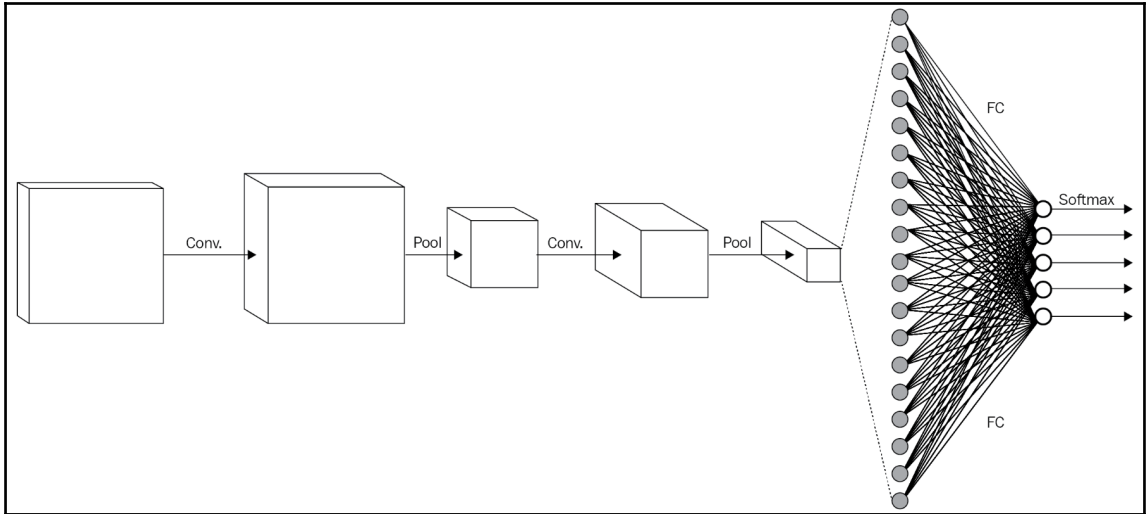
Chapter 1: Data Science - A Birds' Eye View

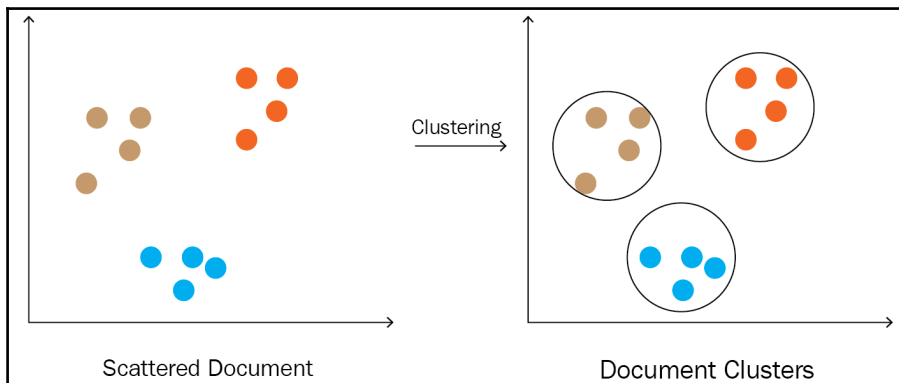
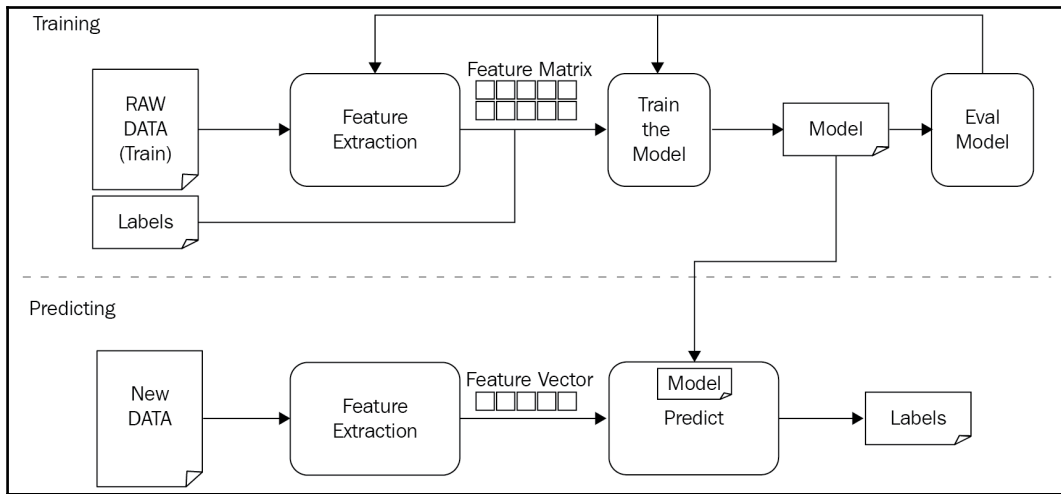
















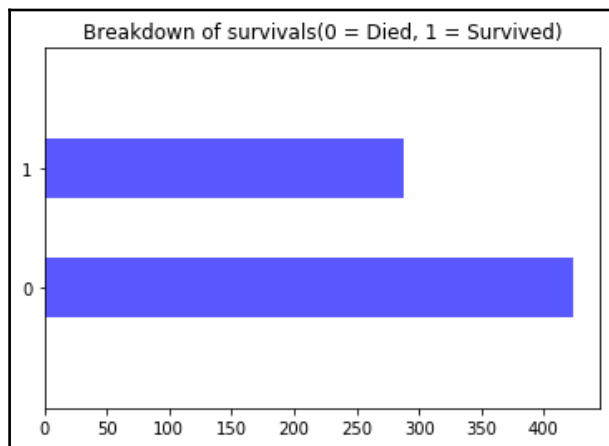
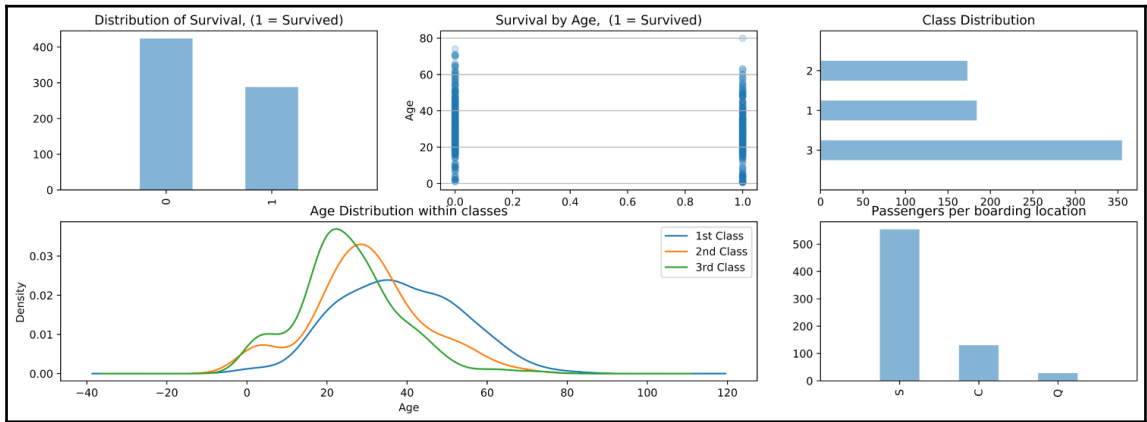


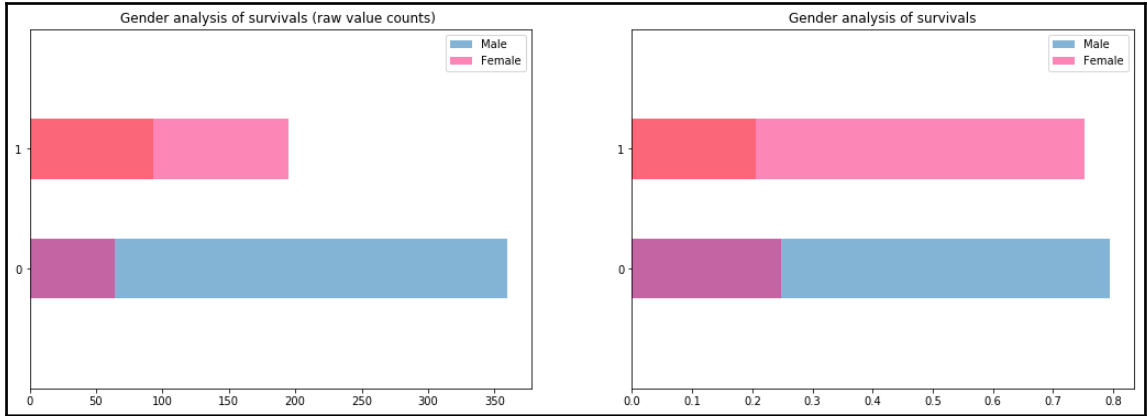
Chapter 2: Data Modeling in Action - The Titanic Example





PassengerId	Survived	Pclass	Name	Sex	Age	SibSp	Parch	Ticket	Fare	Cabin	Embarked
500	0	3	Calic, Mr. Petar	male	17.0	0	0	315086	8.6625	NaN	S
501	0	3	Canavan, Miss. Mary	female	21.0	0	0	364846	7.7500	NaN	Q
502	0	3	O'Sullivan, Miss. Bridget Mary	female	NaN	0	0	330909	7.6292	NaN	Q
503	0	3	Laitinen, Miss. Kristina Sofia	female	37.0	0	0	4135	9.5875	NaN	S
504	1	1	Maioni, Miss. Roberta	female	16.0	0	0	110152	86.5000	B79	S
505	0	1	Penasco y Castellana, Mr. Victor de Satode	male	18.0	1	0	PC 17758	108.9000	C65	C
506	1	2	Quick, Mrs. Frederick Charles (Jane Richards)	female	33.0	0	2	26360	26.0000	NaN	S
507	1	1	Bradley, Mr. George ("George Arthur Brayton")	male	NaN	0	0	111427	26.5500	NaN	S
508	0	3	Olsen, Mr. Henry Margido	male	28.0	0	0	C 4001	22.5250	NaN	S
509	1	3	Lang, Mr. Fang	male	26.0	0	0	1601	56.4958	NaN	S








Logit Regression Results

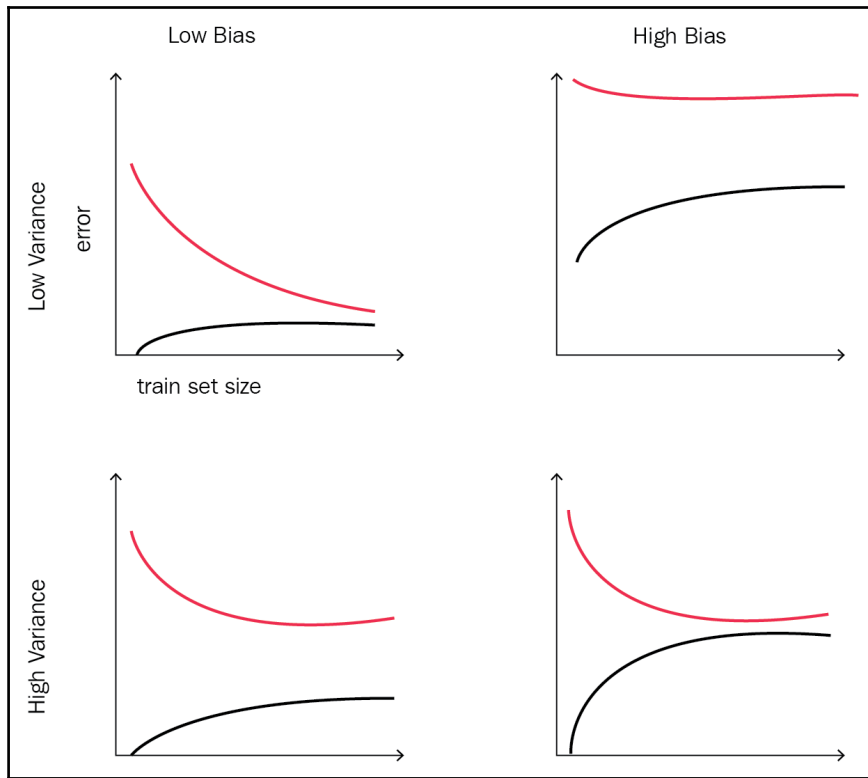
Dep. Variable:	Survived	No. Observations:	712
Model:	Logit	Df Residuals:	704
Method:	MLE	Df Model:	7
Date:	Sun, 20 Dec 2015	Pseudo R-squ.:	0.3414
Time:	11:27:33	Log-Likelihood:	-316.40
converged:	True	LL-Null:	-480.45
		LLR p-value:	5.992e-67

	coef	std err	z	P> z	[95.0% Conf. Int.]
Intercept	4.5423	0.474	9.583	0.000	3.613 5.471
C(Pclass)[T.2]	-1.2673	0.299	-4.245	0.000	-1.852 -0.682
C(Pclass)[T.3]	-2.4966	0.296	-8.422	0.000	-3.078 -1.916
C(Sex)[T.male]	-2.6239	0.218	-12.060	0.000	-3.050 -2.197
C(Embarked)[T.Q]	-0.8351	0.597	-1.398	0.162	-2.006 0.335
C(Embarked)[T.S]	-0.4254	0.271	-1.572	0.116	-0.956 0.105
Age	-0.0436	0.008	-5.264	0.000	-0.060 -0.027
SibSp	-0.3697	0.123	-3.004	0.003	-0.611 -0.129



Chapter 3: Feature Engineering and Model Complexity – The Titanic Example Revisited





Chapter 4: Get Up and Running with TensorFlow

```
ahmed@ahmed-Inspiron-7559:~$ lspci | grep -i nvidia
02:00.0 3D controller: NVIDIA Corporation GM107M [GeForce GTX 960M] (rev a2)
ahmed@ahmed-Inspiron-7559:~$
```

```
ahmed@ahmed-Inspiron-7559:~$ cat /proc/driver/nvidia/version
NVRM version: NVIDIA UNIX x86_64 Kernel Module 375.26 Thu Dec 8 18:36:43 PST
2016
GCC version: gcc version 5.4.0 20160609 (Ubuntu 5.4.0-6ubuntu1~16.04.4)
ahmed@ahmed-Inspiron-7559:~$
```

The screenshot shows the NVIDIA CUDA Toolkit installer website. The main heading is "Select Target Platform". Below this, there is a list of options for Operating System, Architecture, Distribution, Version, and Installer Type. The selected options are Linux, x86_64, Ubuntu, 16.04, and deb [local]. Below the selection options, there is a section titled "Download Installers for Linux Ubuntu 16.04 x86_64". This section contains the text: "The base installer is available for download below. There is 1 patch available. This patch requires the base installer to be installed first." Below this text, there is a button labeled "Base Installer" and a "Download (1.9 GB)" button. Below the buttons, there are "Installation Instructions:" which are:

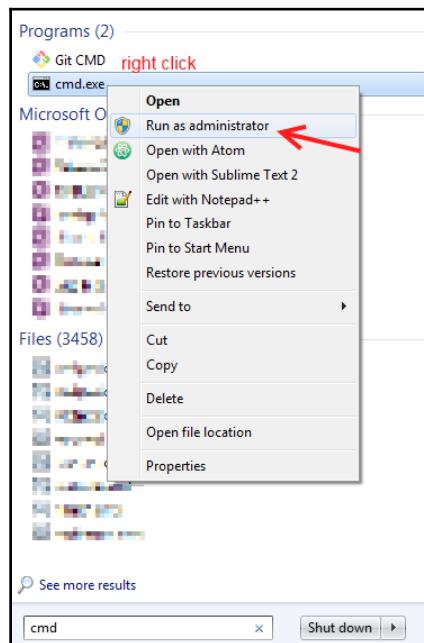
1. ``sudo dpkg -i cuda-repo-ubuntu1604-8-0-local-ga2_8.0.61-1_amd64.deb``
2. ``sudo apt-get update``
3. ``sudo apt-get install cuda``




```
ahmed@ahmed-Inspiron-7559:~$ nvcc -V
nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2016 NVIDIA Corporation
Built on Tue_Jan_10_13:22:03_CST_2017
Cuda compilation tools, release 8.0, V8.0.61
ahmed@ahmed-Inspiron-7559:~$
```


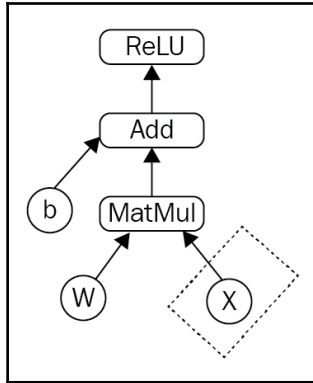
```
ahmed@ahmed-Inspiron-7559:~$ python3
Python 3.6.0 |Anaconda 4.3.0 (64-bit)| (default, Dec 23 2016, 12:22:00)
[GCC 4.4.7 20120313 (Red Hat 4.4.7-1)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import tensorflow as tf
I tensorflow/stream_executor/dso_loader.cc:128] successfully opened CUDA library libcublas.so locally
I tensorflow/stream_executor/dso_loader.cc:128] successfully opened CUDA library libcudnn.so locally
I tensorflow/stream_executor/dso_loader.cc:128] successfully opened CUDA library libcufft.so locally
I tensorflow/stream_executor/dso_loader.cc:128] successfully opened CUDA library libcuda.so.1 locally
I tensorflow/stream_executor/dso_loader.cc:128] successfully opened CUDA library libcurand.so locally
>>> a = tf.constant(5)
>>> b = tf.constant(6)
>>> sess = tf.Session()
I tensorflow/stream_executor/cuda/cuda_gpu_executor.cc:937] successful NUMA node read from SysFS had
negative value (-1), but there must be at least one NUMA node, so returning NUMA node zero
I tensorflow/core/common_runtime/gpu/gpu_device.cc:885] Found device 0 with properties:
name: GeForce GTX 960M
major: 5 minor: 0 memoryClockRate (GHz) 1.176
pciBusID 0000:02:00.0
Total memory: 3.95GiB
Free memory: 3.43GiB
I tensorflow/core/common_runtime/gpu/gpu_device.cc:906] DMA: 0
I tensorflow/core/common_runtime/gpu/gpu_device.cc:916] 0: Y
I tensorflow/core/common_runtime/gpu/gpu_device.cc:975] Creating TensorFlow device (/gpu:0) -> (devic
e: 0, name: GeForce GTX 960M, pci bus id: 0000:02:00.0)
>>> sess.run(a+b)
11
>>> sess.close()
>>>
```



```
ahmed@ahmed-Inspiron-7559:~$ python3
Python 3.6.0 [Anaconda 4.3.0 (64-bit)] (default, Dec 23 2016, 12:22:00)
[GCC 4.4.7 20120313 (Red Hat 4.4.7-1)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import tensorflow as tf
>>> a = tf.constant(5)
>>> b = tf.constant(6)
>>> sess = tf.Session()
2017-09-11 15:39:50.792072: W tensorflow/core/platform/cpu_feature_guard.cc:45] The TensorFlow library wasn't compiled to use SSE4.1 instructions, but these are available on your machine and could speed up CPU computations.
2017-09-11 15:39:50.792125: W tensorflow/core/platform/cpu_feature_guard.cc:45] The TensorFlow library wasn't compiled to use SSE4.2 instructions, but these are available on your machine and could speed up CPU computations.
2017-09-11 15:39:50.792143: W tensorflow/core/platform/cpu_feature_guard.cc:45] The TensorFlow library wasn't compiled to use AVX instructions, but these are available on your machine and could speed up CPU computations.
2017-09-11 15:39:50.792158: W tensorflow/core/platform/cpu_feature_guard.cc:45] The TensorFlow library wasn't compiled to use AVX2 instructions, but these are available on your machine and could speed up CPU computations.
2017-09-11 15:39:50.792177: W tensorflow/core/platform/cpu_feature_guard.cc:45] The TensorFlow library wasn't compiled to use FMA instructions, but these are available on your machine and could speed up CPU computations.
>>> sess.run(a+b)
11
>>> sess.close()
>>> █
```







```

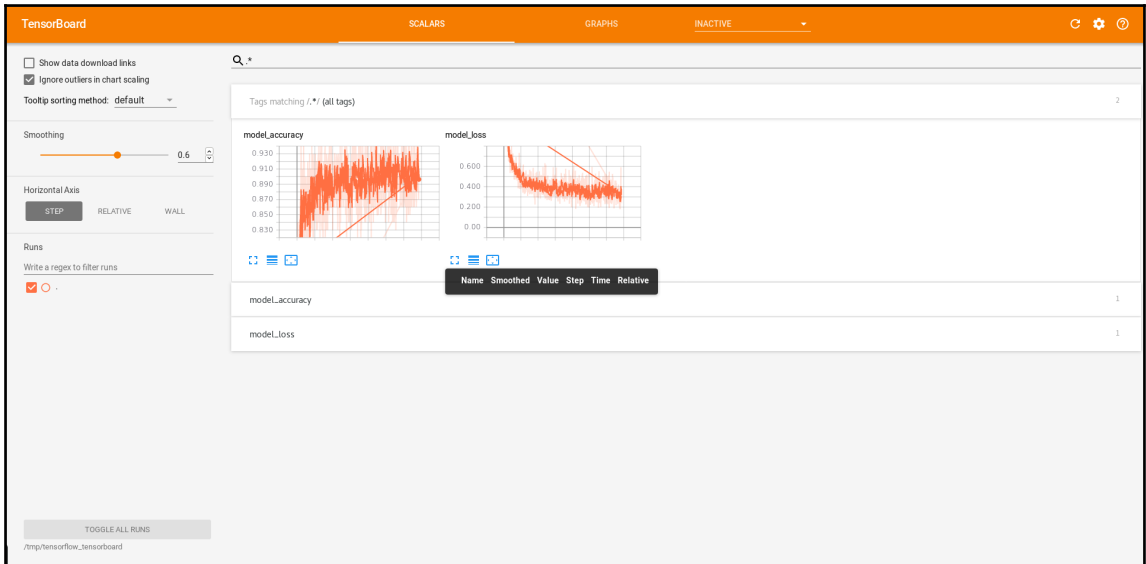
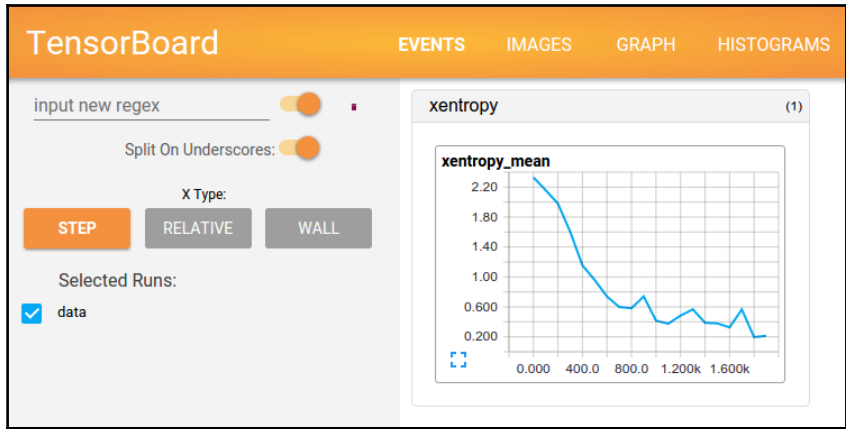
ahmed@ahmed-Inspiron-7559:~$ python3
Python 3.6.0 |Anaconda 4.3.0 (64-bit)| (default, Dec 23 2016, 12:22:00)
[GCC 4.4.7 20120313 (Red Hat 4.4.7-1)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> # import TensorFlow package
... import tensorflow as tf
>>> # build a TensorFlow variable b taking in initial zeros of size 100
... # ( a vector of 100 values)
... b = tf.Variable(tf.zeros((100,)))
>>> # TensorFlow variable uniformly distributed values between -1 and 1
... # of shape 784 by 100
... W = tf.Variable(tf.random_uniform((784, 100),-1,1))
>>> # TensorFlow placeholder for our input data that doesn't take in
... # any initial values, it just takes a data type 32 bit floats as
... # well as its shape
... x = tf.placeholder(tf.float32, (100, 784))
>>> # express h as Tensorflow ReLU of the TensorFlow matrix
... #Multiplication of x and W and we add b
... h = tf.nn.relu(tf.matmul(x,W) + b )
>>> print(W)
<tf.Variable 'Variable_1:0' shape=(784, 100) dtype=float32_ref>
>>> print(b)
<tf.Variable 'Variable:0' shape=(100,) dtype=float32_ref>
>>> █

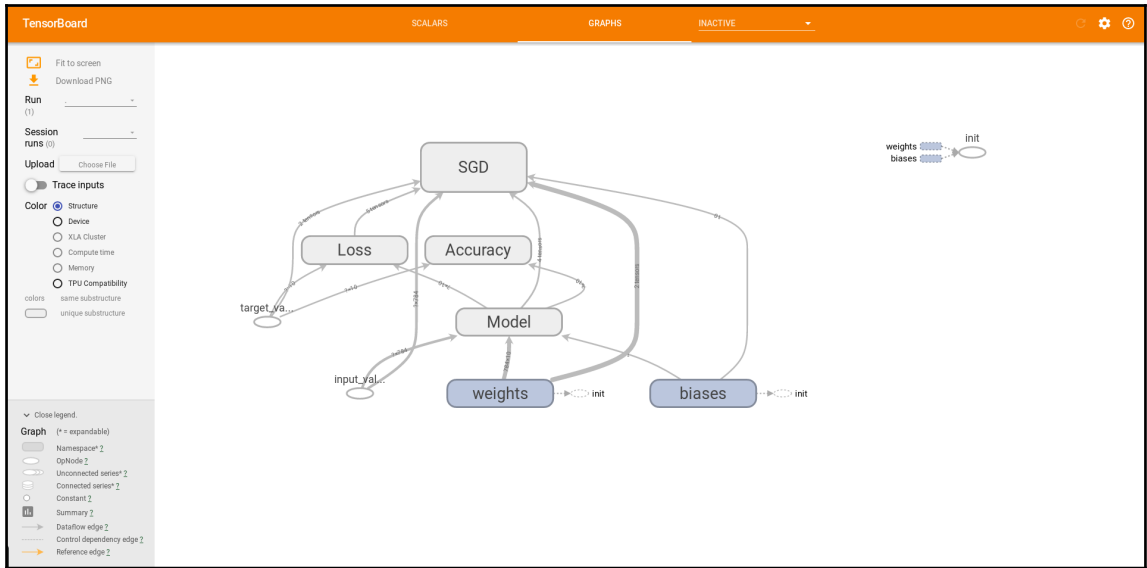
```

```




>>> # importing the numpy package for generating random variables for
... # our placeholder x
... import numpy as np
>>> # build a TensorFlow session object which takes a default execution
... # environment which will be most likely a CPU
... sess = tf.Session()
>>> # calling the run function of the sess object to initialize all the
... # variables.
... sess.run(tf.global_variables_initializer())
>>> # calling the run function on the node that we are interested in,
... # the h, and we feed in our second argument which is a dictionary
... # for our placeholder x with the values that we are interested in.
... sess.run(h, {x: np.random.random((100,784))})
array([[ 4.95583916,  0.13156724,  0.          , ...,  0.          ,
         0.          ,  0.          ],
       [ 0.          ,  0.          ,  0.          , ...,  0.          ,
         0.          ,  0.          ],
       [ 0.          ,  0.          ,  0.          , ...,  0.          ,
         0.          ,  0.          ],
       ...,
       [ 2.81067681,  8.28696823,  0.          , ...,  0.          ,
         0.55022001,  0.          ],
       [ 0.          ,  6.67730427,  0.          , ...,  0.          ,
         4.28411198,  3.10559845],
       [ 1.92718267,  0.          ,  0.          , ...,  0.          ,
         0.          ,  0.          ]], dtype=float32)
>>> █




```

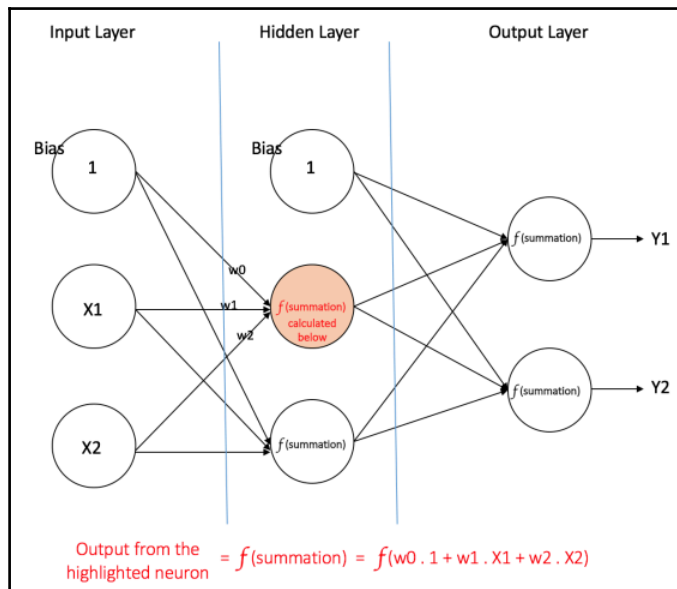









Chapter 5: TensorFlow in Action - Some Basic Examples












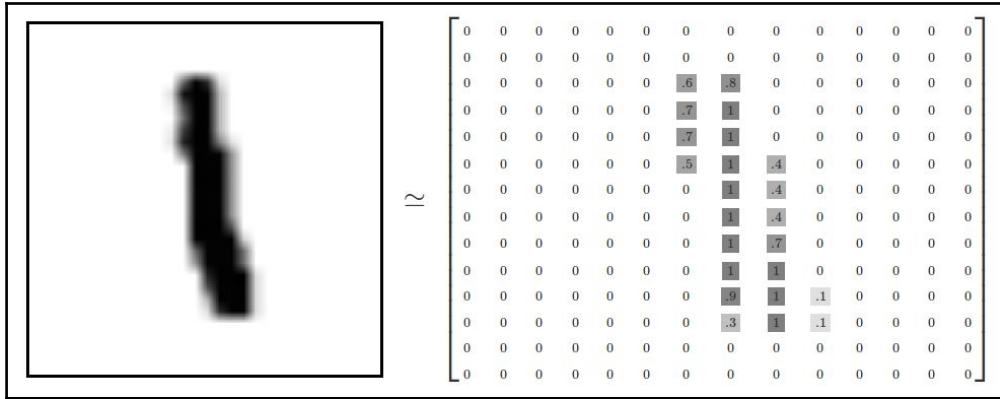








Chapter 6: Deep Feed-forward Neural Networks - Implementing Digit Classification














$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \text{softmax} \begin{bmatrix} W_{1,1}x_1 + W_{1,2}x_2 + W_{1,3}x_3 + b_1 \\ W_{2,1}x_1 + W_{2,2}x_2 + W_{2,3}x_3 + b_2 \\ W_{3,1}x_1 + W_{3,2}x_2 + W_{3,3}x_3 + b_3 \end{bmatrix}$$

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \text{softmax} \left(\begin{bmatrix} W_{1,1} & W_{1,2} & W_{1,3} \\ W_{2,1} & W_{2,2} & W_{2,3} \\ W_{3,1} & W_{3,2} & W_{3,3} \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} \right)$$

Label	0	1	2	3	4	5	6	7	8	9
Array	[0,	0,	0,	0,	0,	0,	0,	1,	0,	0]







Chapter 7: Introduction to Convolutional Neural Networks

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

1	0	1
0	1	0
1	0	1

1 _{x1}	1 _{x0}	1 _{x1}	0	0
0 _{x0}	1 _{x1}	1 _{x0}	1	0
0 _{x1}	0 _{x0}	1 _{x1}	1	1
0	0	1	1	0
0	1	1	0	0

Image

4		





Convolved Feature





1	1 _{x1}	1 _{x0}	0 _{x1}	0
0	1 _{x0}	1 _{x1}	1 _{x0}	0
0	0 _{x1}	1 _{x0}	1 _{x1}	1
0	0	1	1	0
0	1	1	0	0



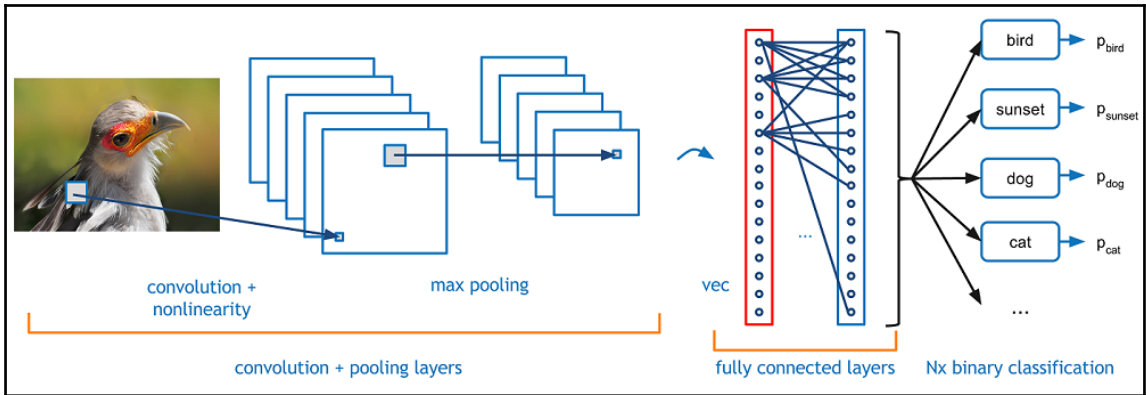
Image



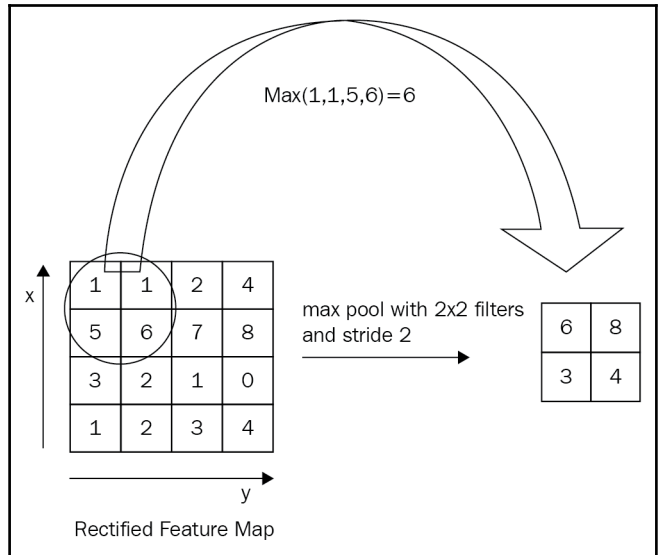
4	3	


Convolved Feature




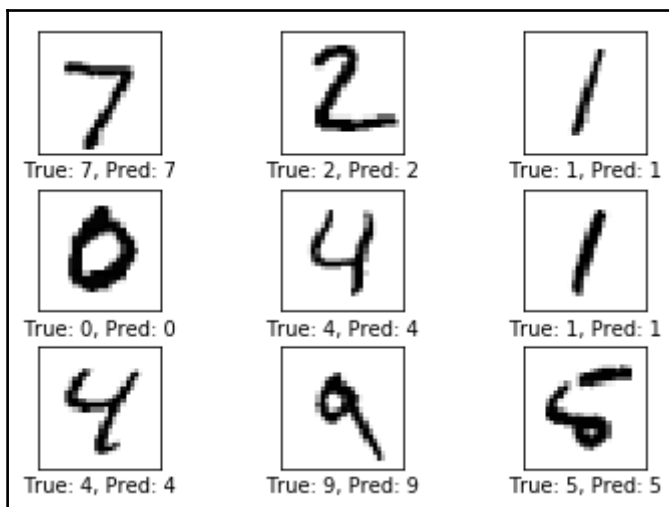








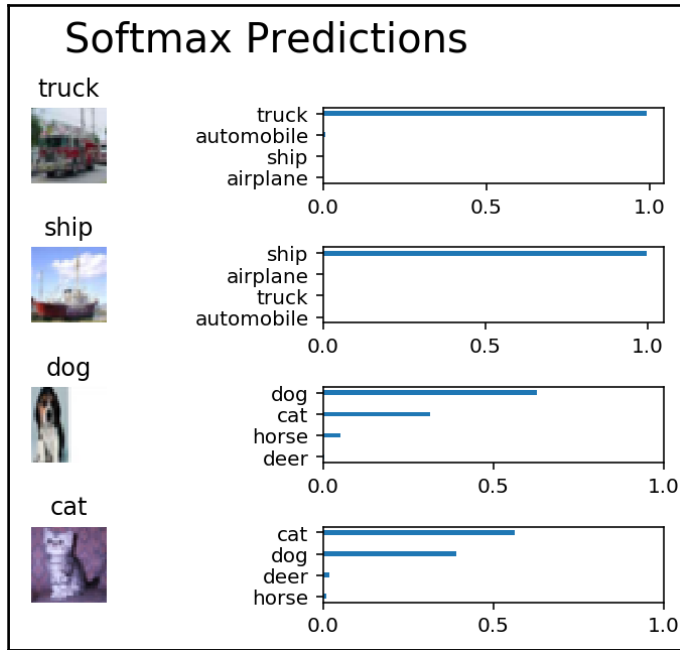


True: 7	True: 2	True: 1
True: 0	True: 4	True: 1
True: 4	True: 9	True: 5

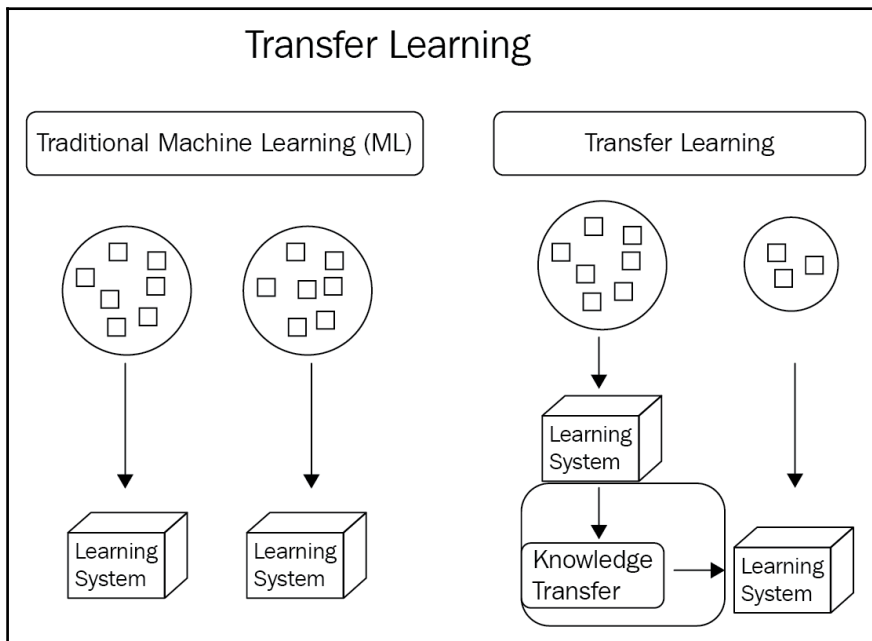
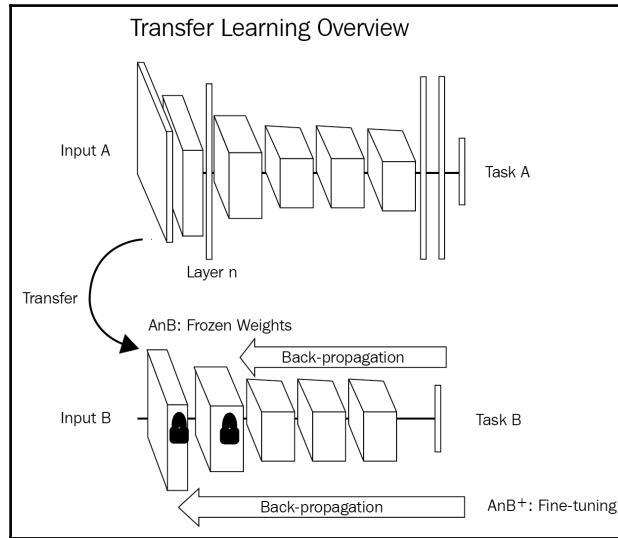


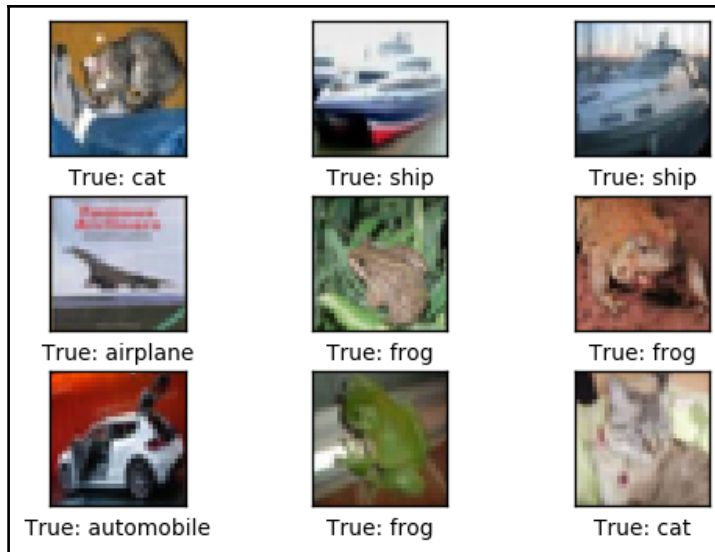
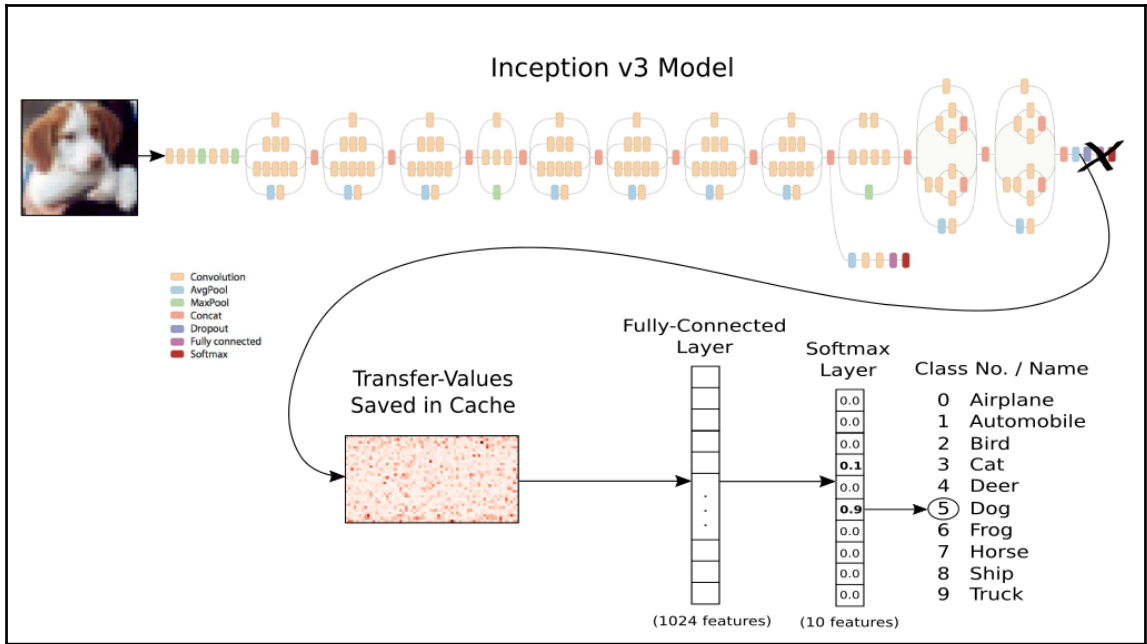
Chapter 8: Object Detection – CIFAR-10 Example











Chapter 9: Object Detection – Transfer Learning with CNNs



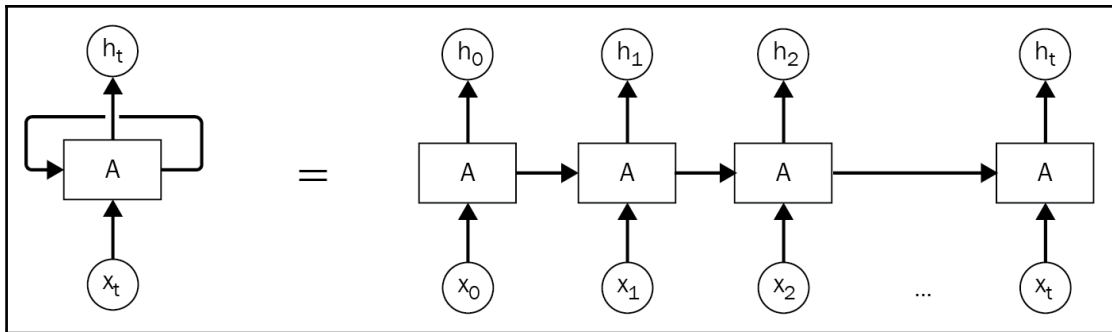

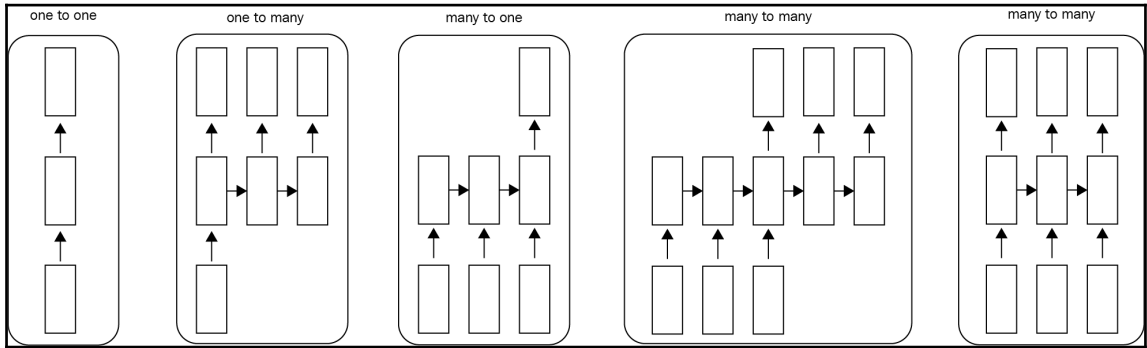


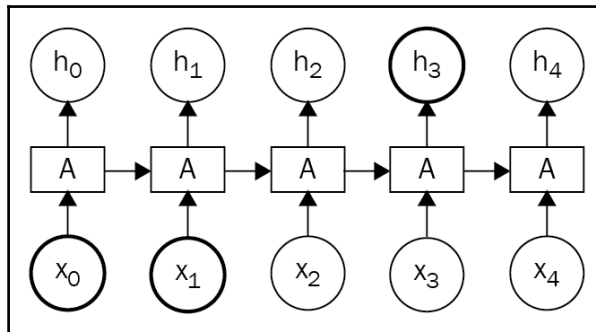

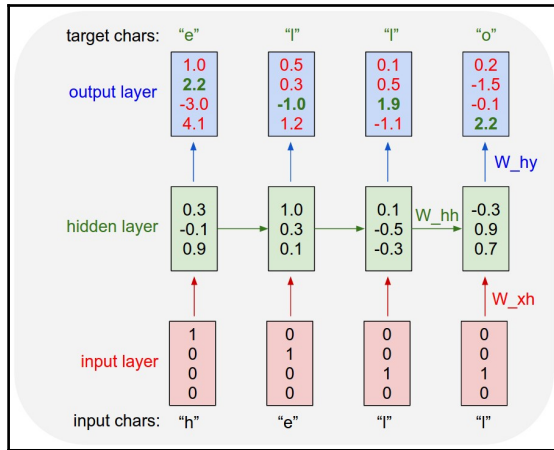


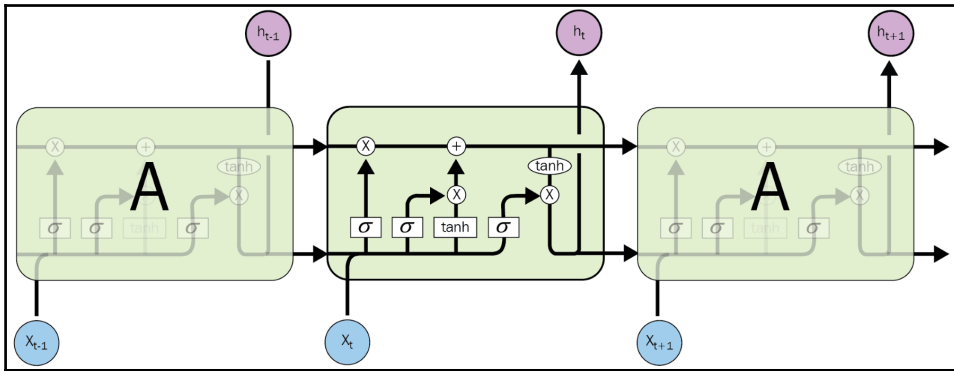

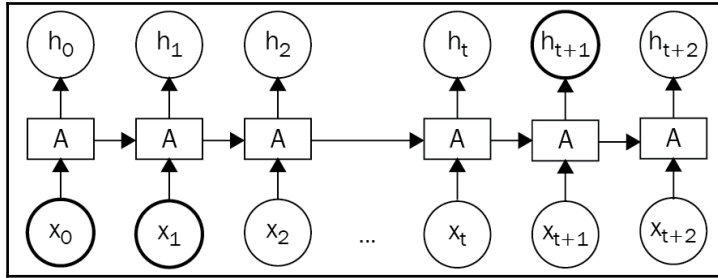



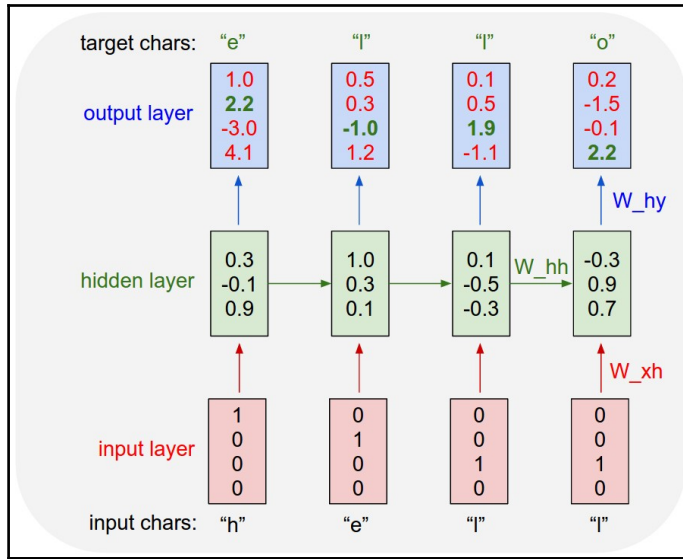


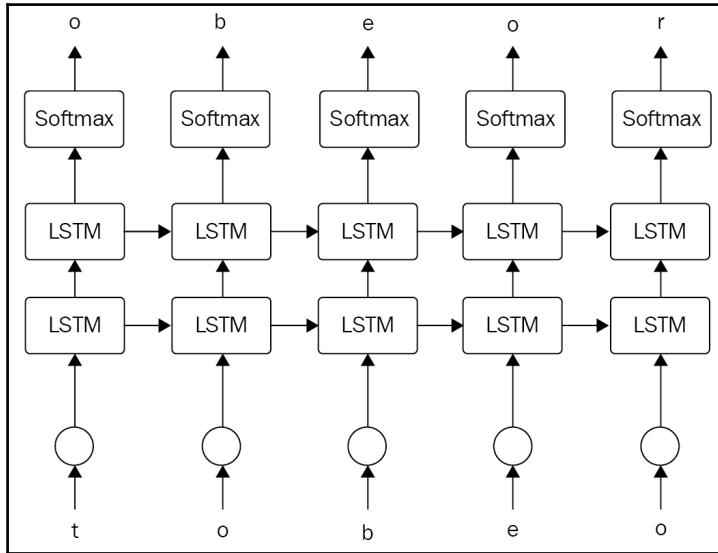

Chapter 10: Recurrent-Type Neural Networks - Language Modeling



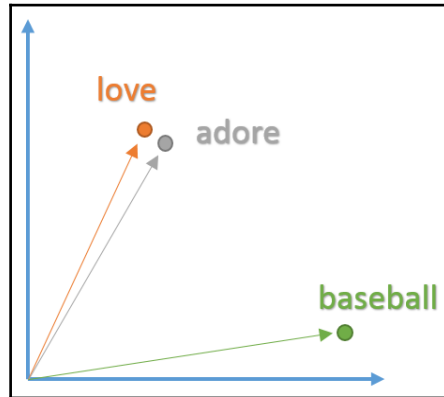








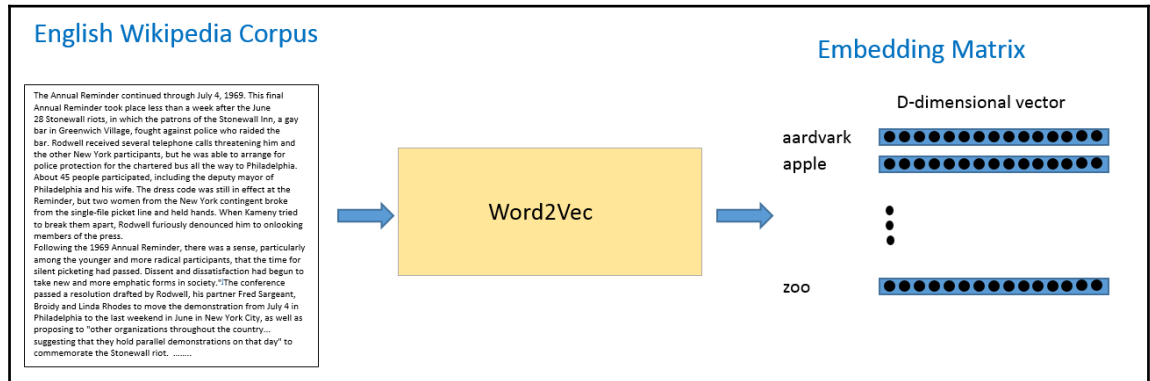
Chapter 11: Representation Learning - Implementing Word Embeddings

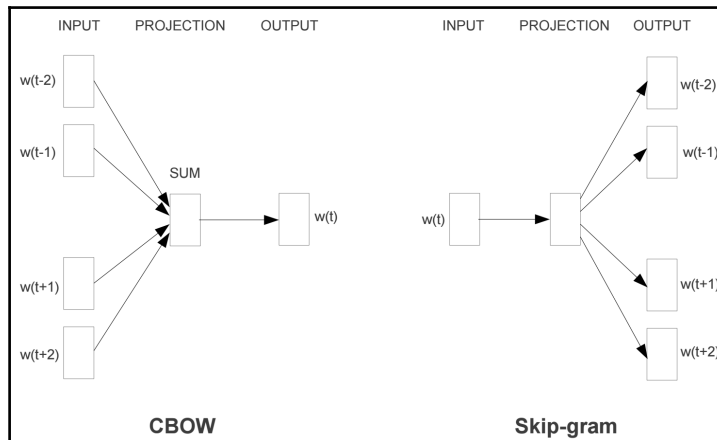
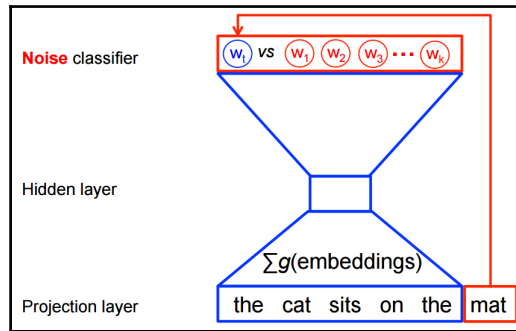
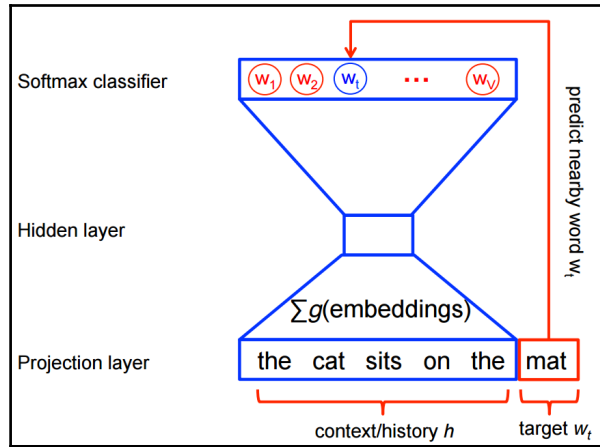


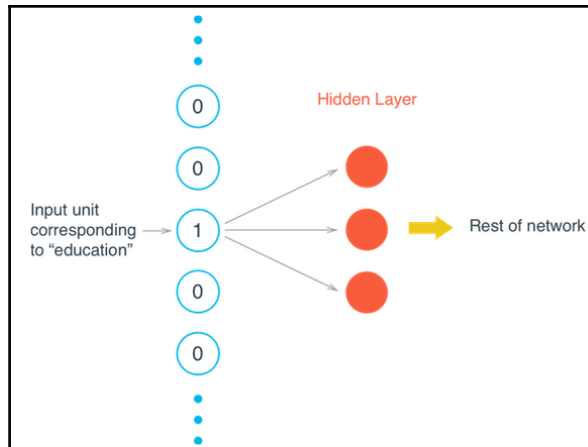

I love taking long walks on the beach.
 My friends told me that they love popcorn.

•
 •
 •

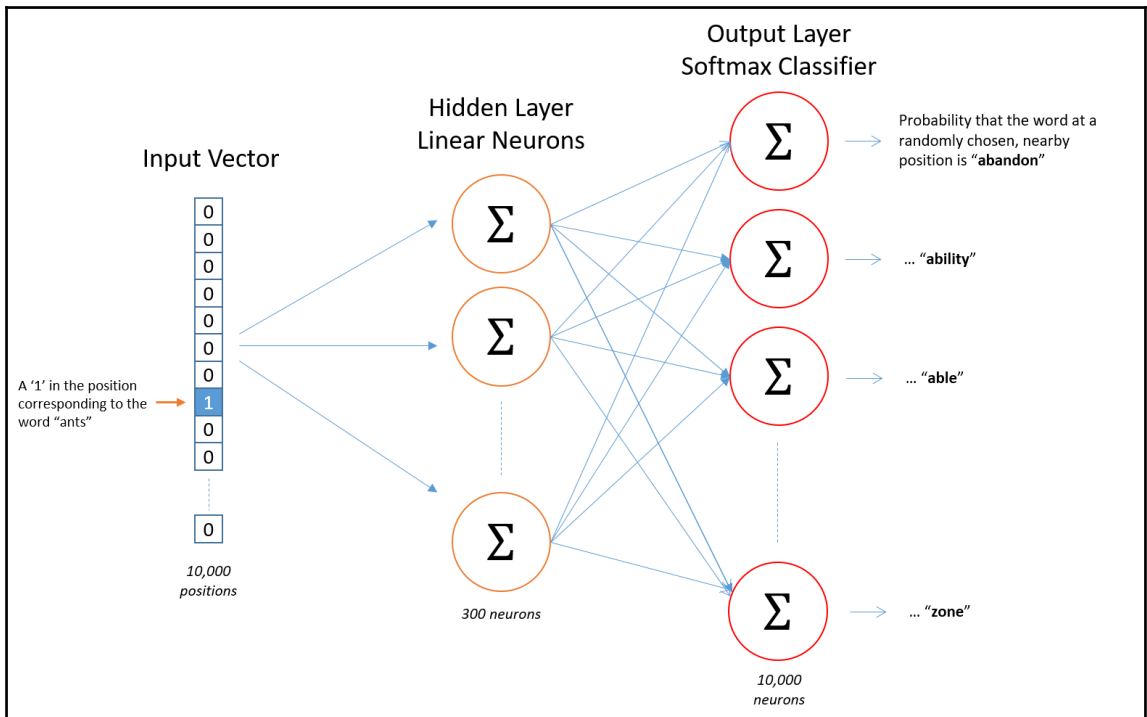

The relatives adore the baby's cute face.
 I adore his sense of humor.







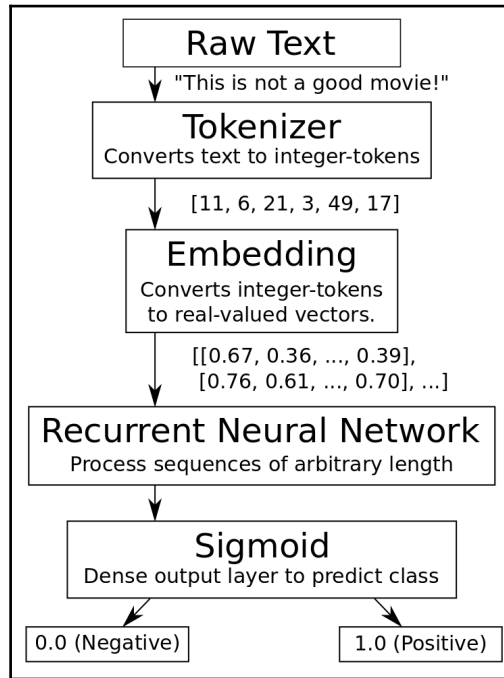


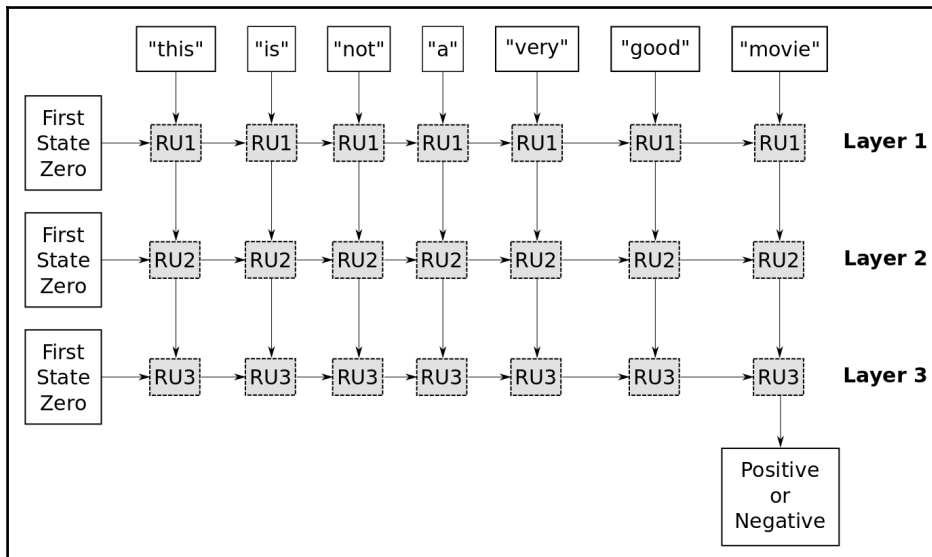
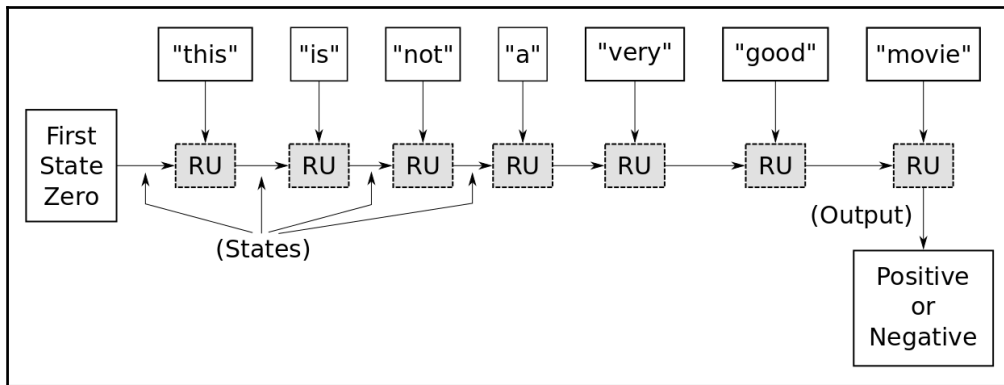
$$\begin{bmatrix} 0 & 0 & 0 & 1 & 0 \end{bmatrix} \times \begin{bmatrix} 8 & 2 & 1 & 9 \\ 6 & 5 & 4 & 0 \\ 7 & 1 & 6 & 2 \\ 1 & 3 & 5 & 8 \\ 0 & 4 & 9 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 3 & 5 & 8 \end{bmatrix}$$









Chapter 12: Neural Sentiment Analysis






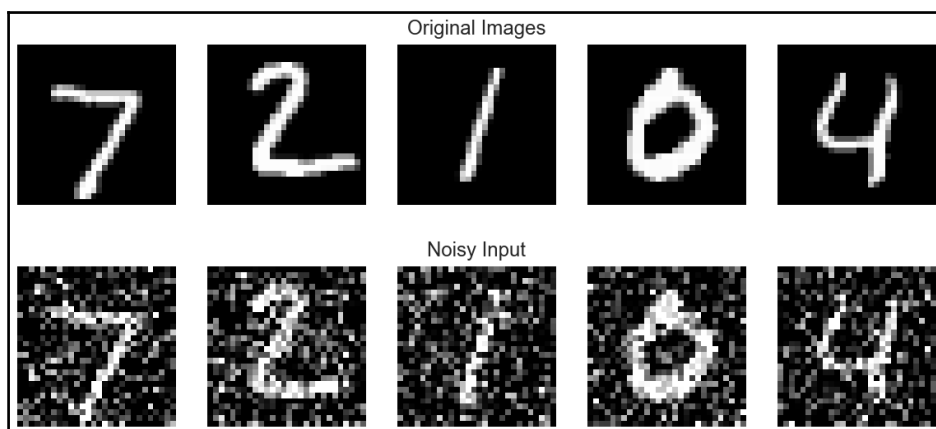







Chapter 13: Autoencoders – Feature Extraction and Denoising





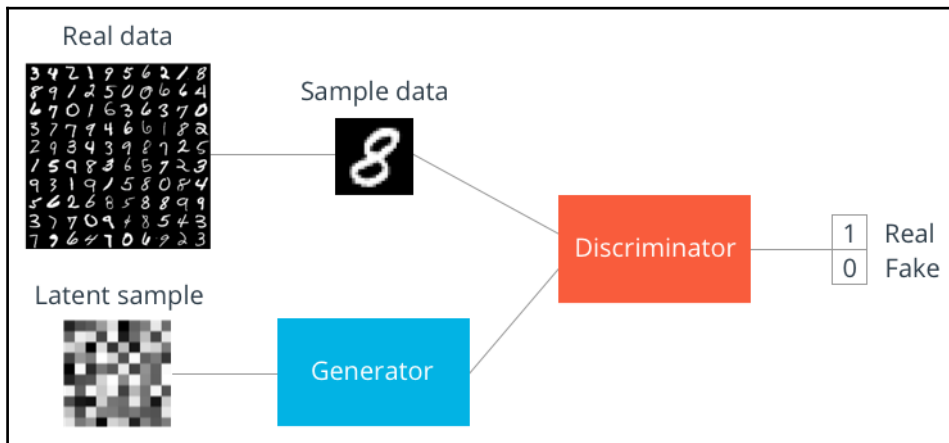




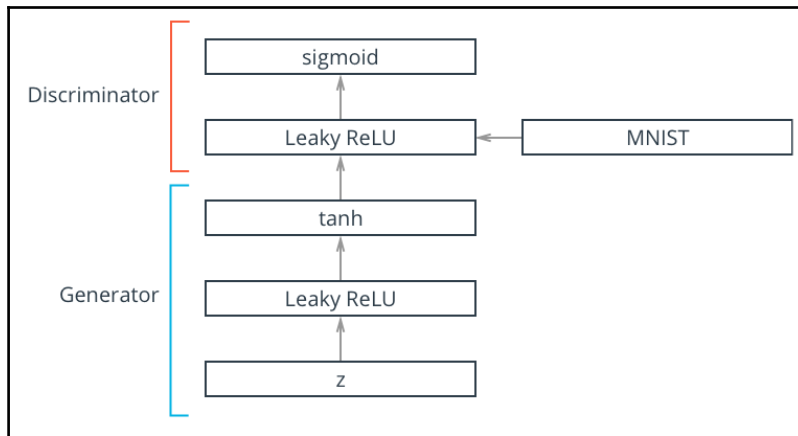


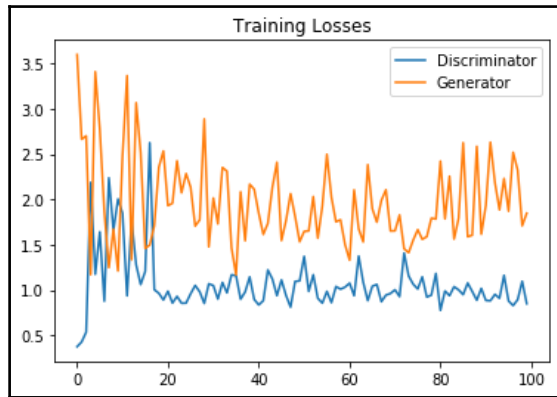








Chapter 14: Generative Adversarial Networks











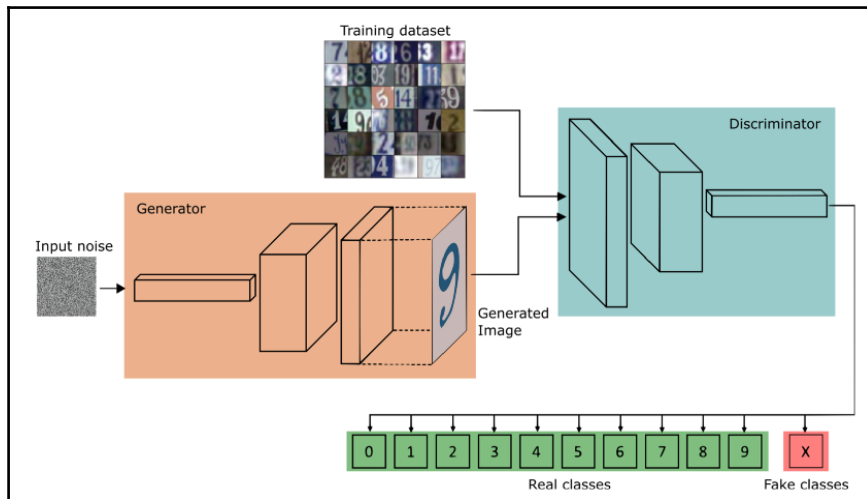







Chapter 15: Face Generation and Handling Missing Labels











Index