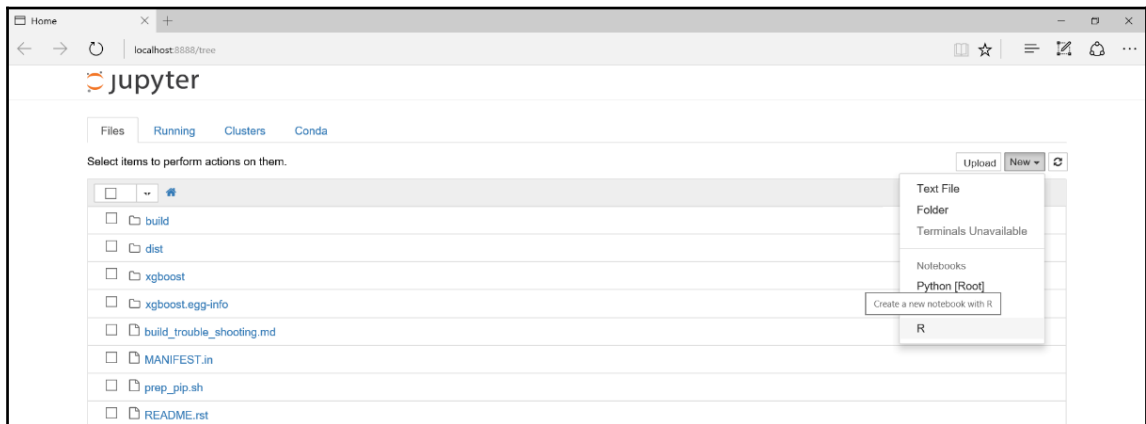
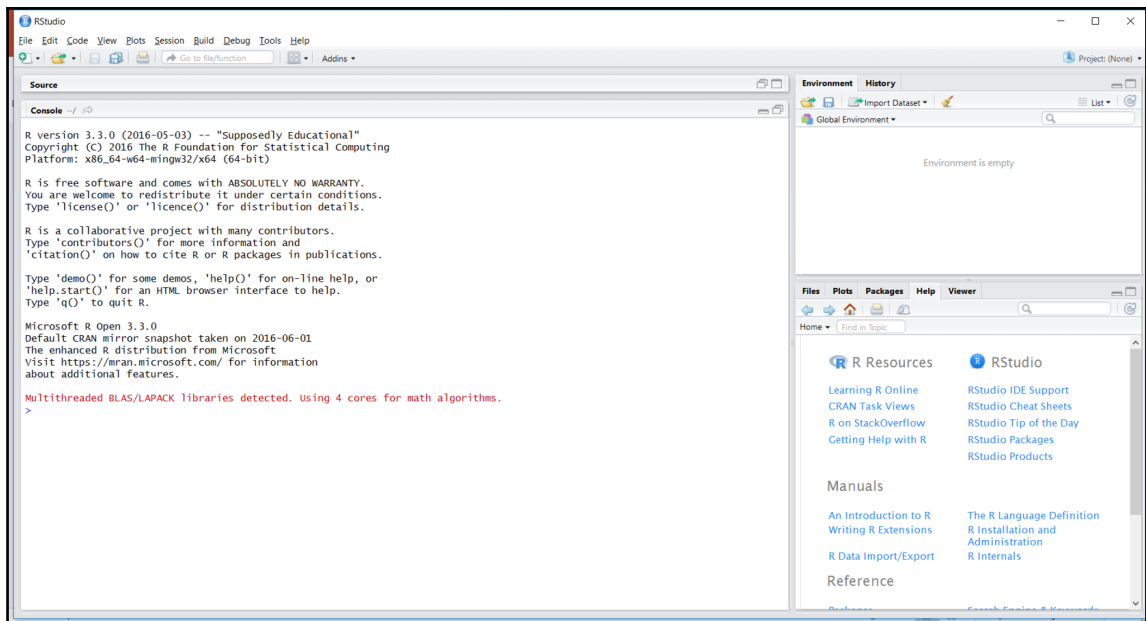


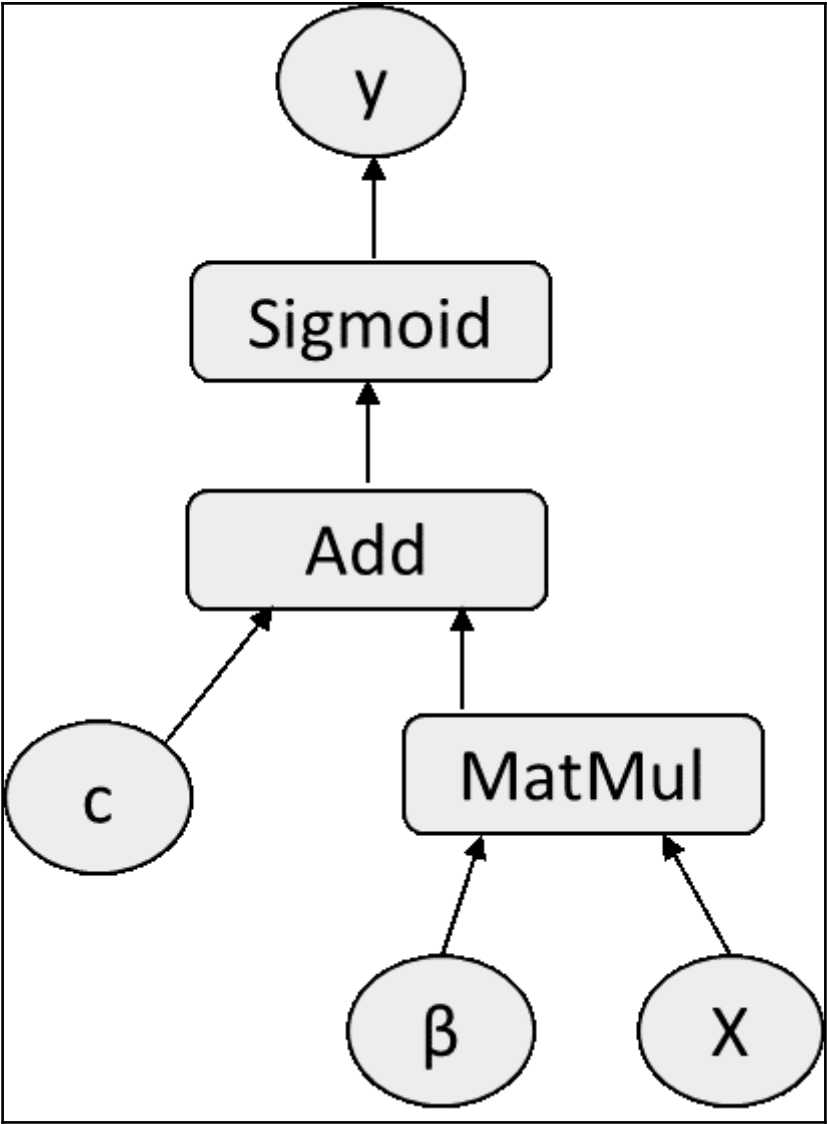
1

Graphics bundle

Chapter 1: Getting Started







```
In [17]: # Loading and setting placeholder in tensorflow
Sys.setenv(TENSORFLOW_PYTHON="/usr/bin/python")
require(tensorflow)
x <- tf$placeholder(tf$float32, shape(NULL, 11L))
```

Error: Python module tensorflow was not found.

Detected Python configuration:

```
python:      /usr/bin/python
libpython:   /usr/lib/python2.7/config-x86_64-linux-gnu/libpython2.7.so
pythonhome:  /usr:/usr
version:     2.7.12 (default, Nov 19 2016, 06:48:10) [GCC 5.4.0 20160609]
numpy:       /usr/lib/python2.7/dist-packages/numpy
numpy_version: 1.11.0
tensorflow:  [NOT FOUND]
```

python versions found:

```
/usr/bin/python
/usr/bin/python3
```

Traceback:

```
1. tf$placeholder
2. `$.python.builtin.module`(tf, placeholder)
3. py_resolve_module_proxy(x)
4. stop(message, call. = FALSE)
```

The screenshot shows the RStudio interface. The console pane on the left displays the output of the `h2o` package initialization. The Environment pane on the right shows the `localH2O` object as a formal class `H2OConnection`.

Console output:

```
The following objects are masked from 'package:base':
  %%, %in%, &&, ||, apply, as.factor, as.numeric, colnames, colnames<-, ifelse,
  is.character, is.factor, is.numeric, log, log10, log1p, log2, round, signif, trunc

> localH2O = h2o.init()

H2O is not running yet, starting it now...

Note: In case of errors look at the following log files:
C:\Users\prakash\AppData\Local\Temp\RtmpCofpin\h2o_prakash_started_from_r.out
C:\Users\prakash\AppData\Local\Temp\RtmpCofpin\h2o_prakash_started_from_r.err

Java version "1.8.0_111"
Java(TM) SE Runtime Environment (build 1.8.0_111-b14)
Java HotSpot(TM) 64-Bit Server VM (build 25.111-b14, mixed mode)

Starting H2O JVM and connecting: Connection successful!

R is connected to the H2O cluster:
H2O cluster uptime:      1 seconds 536 milliseconds
H2O cluster version:    3.8.1.3
H2O cluster name:       H2O_started_from_R_prakash_lvs514
H2O cluster total nodes: 1
H2O cluster total memory: 14.22 GB
H2O cluster total cores: 8
H2O cluster allowed cores: 2
H2O cluster healthy:    TRUE
H2O Connection ip:      localhost
H2O Connection port:    54321
H2O Connection proxy:   NA
R Version:              R version 3.3.0 (2016-05-03)

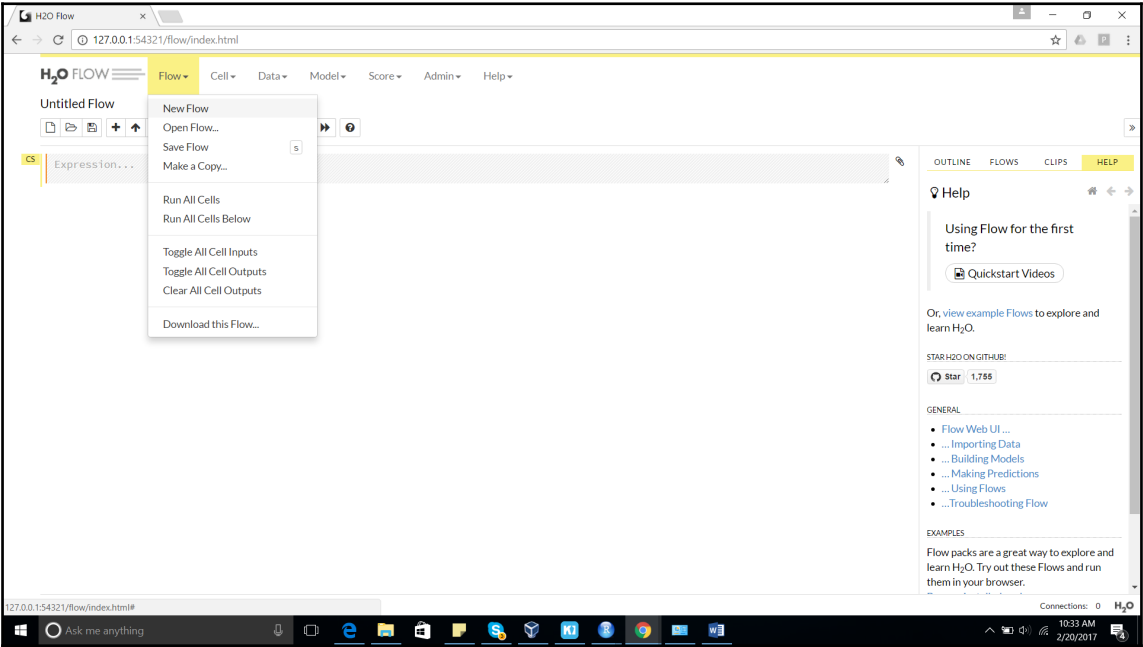
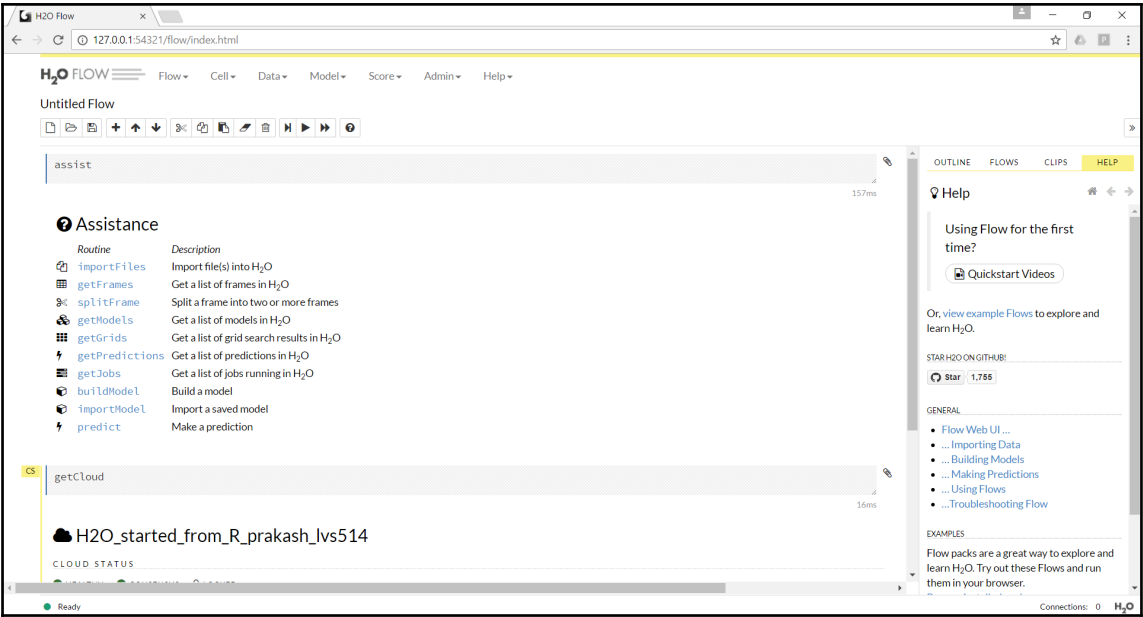
Note: As started, H2O is limited to the CRAN default of 2 CPUS.
Shut down and restart H2O as shown below to use all your CPUS.
> h2o.shutdown()
> h2o.init(nthreads = -1)
```

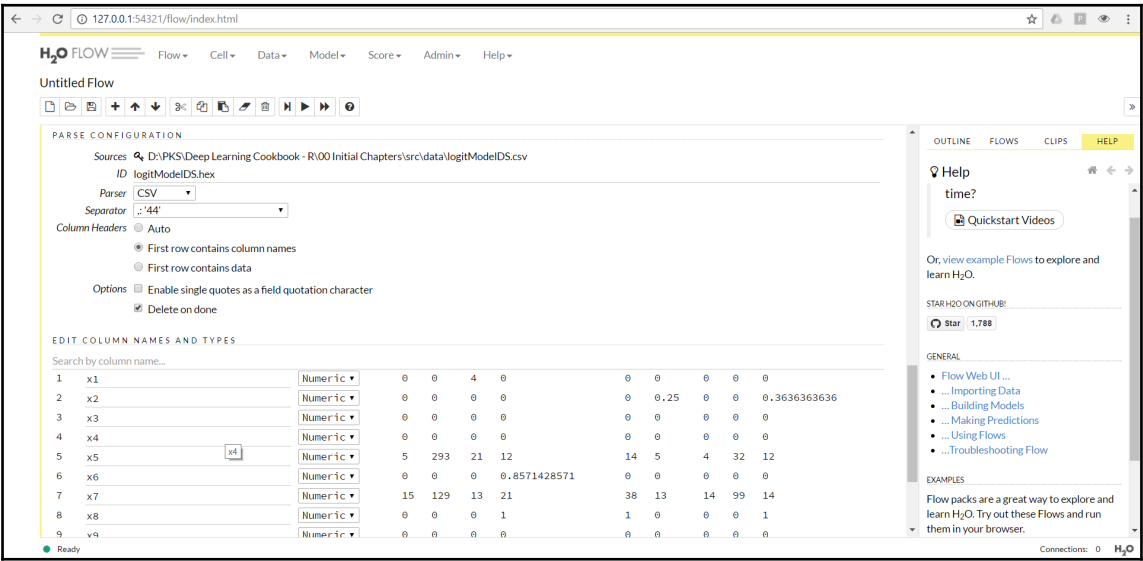
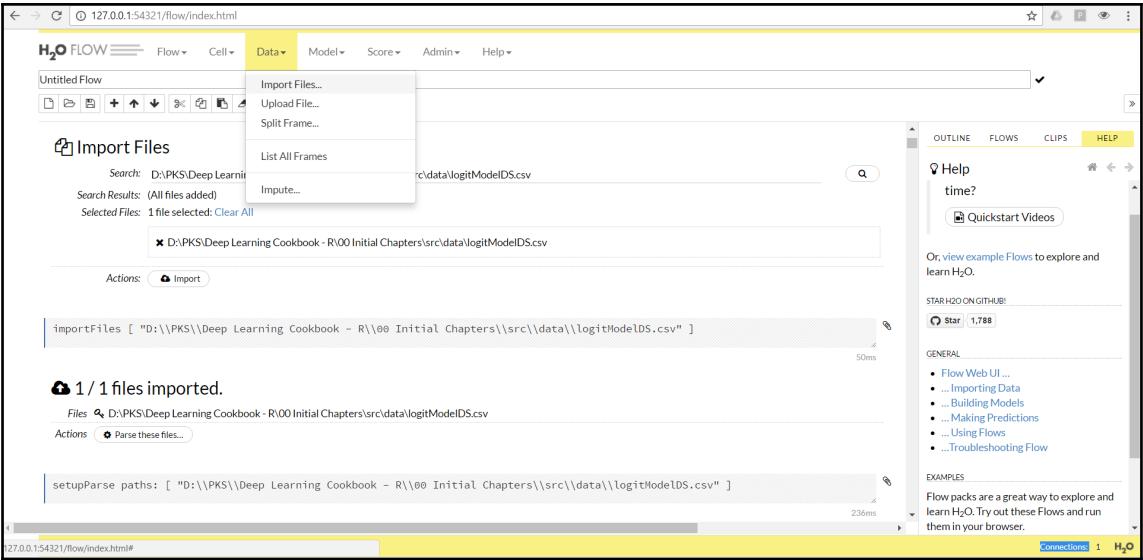
Environment pane:

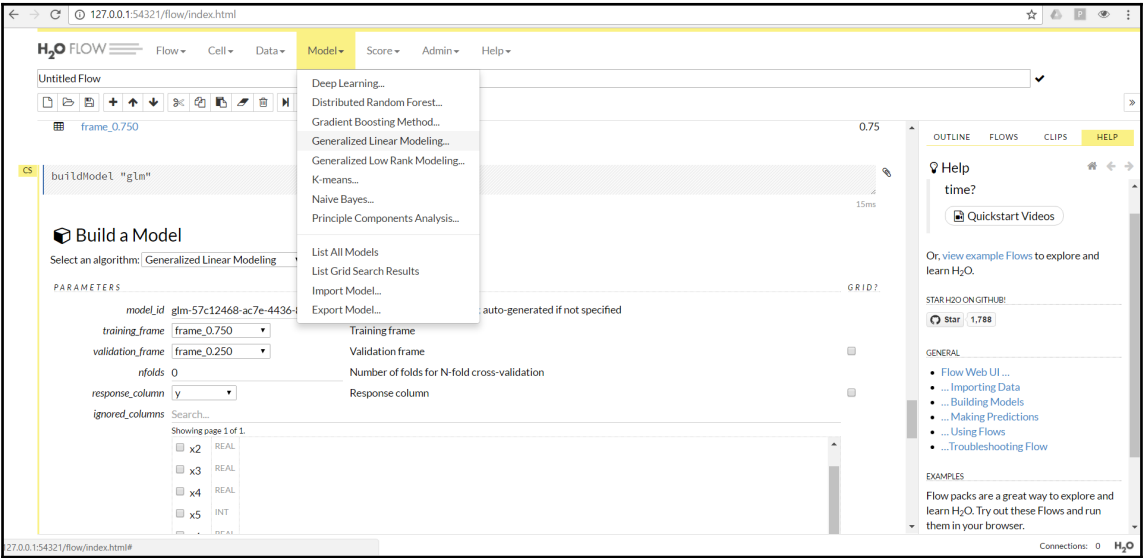
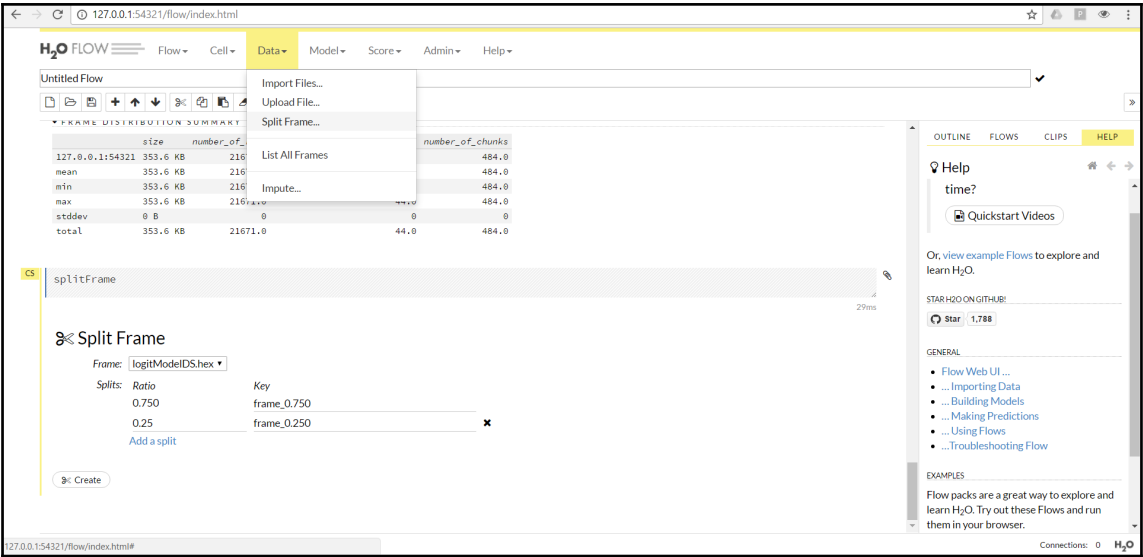
Values
localH2O




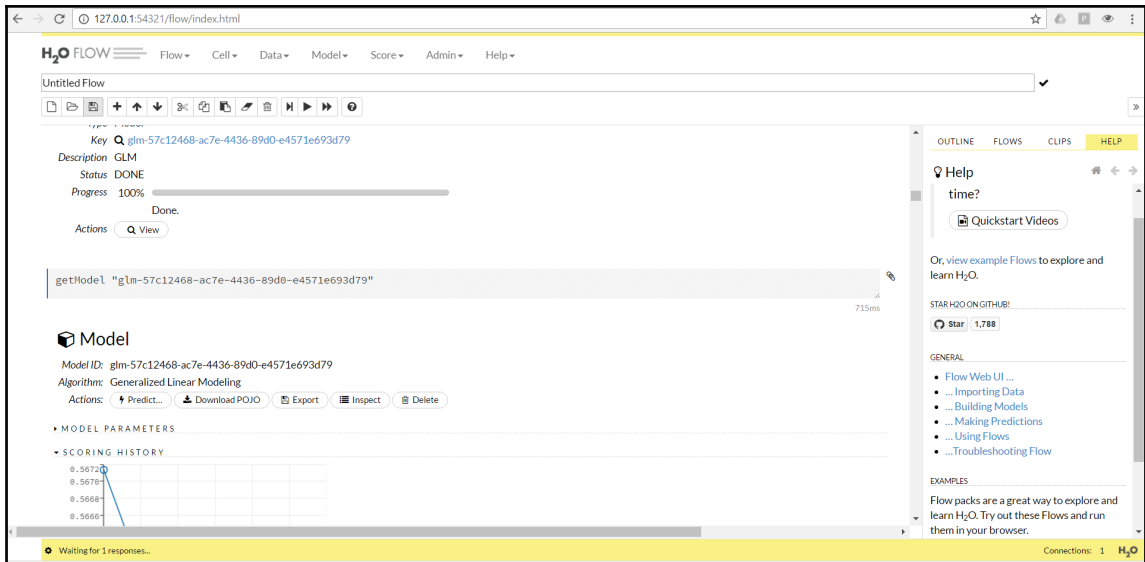
Files pane:

Files
h2o.jar








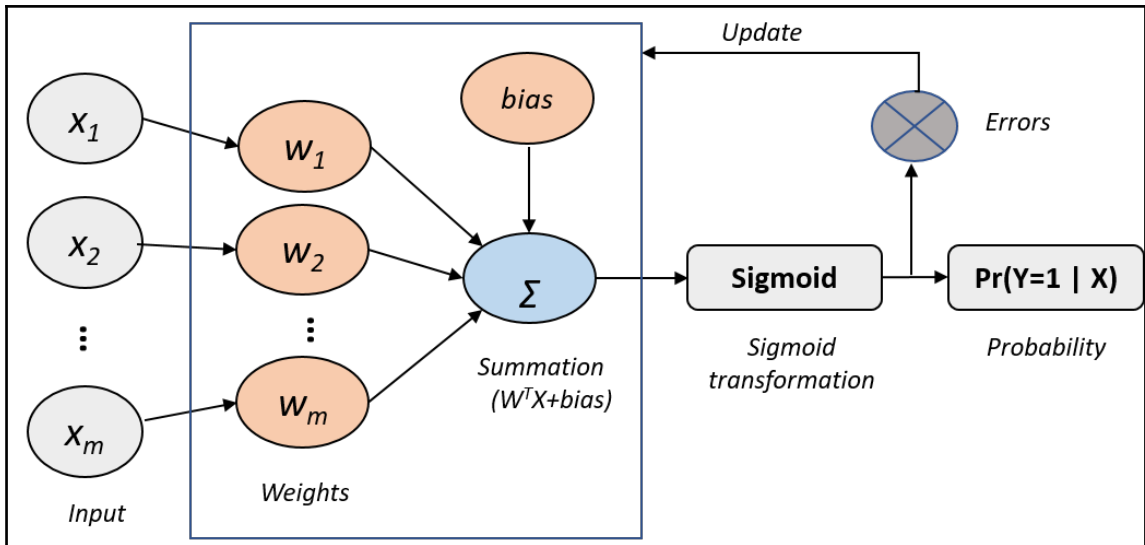


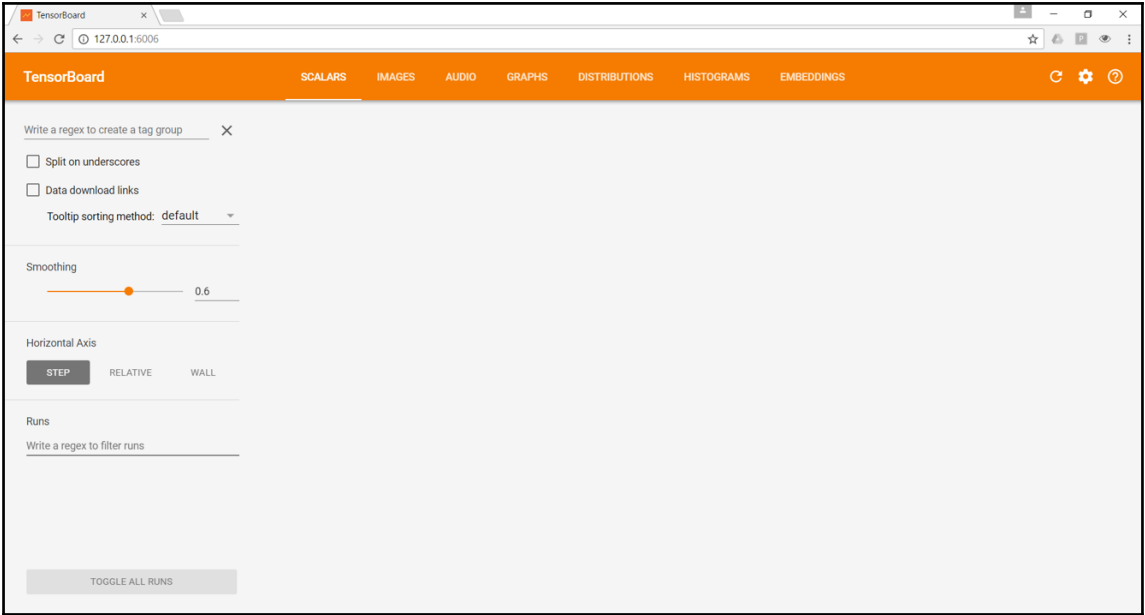




Chapter 2: Deep Learning with R

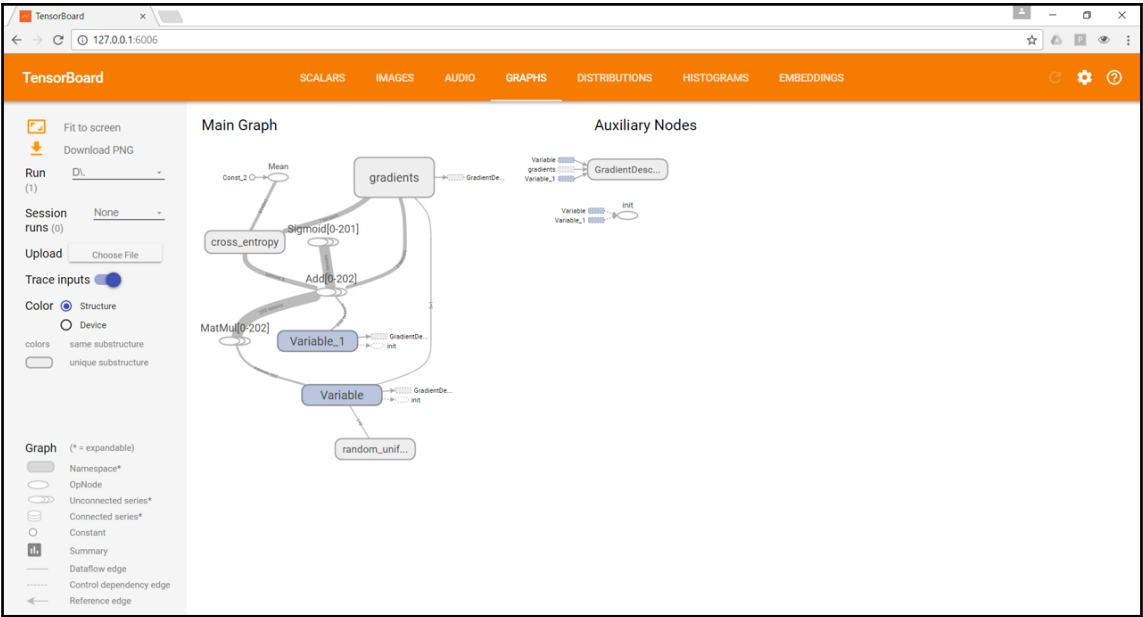

























Chapter 3: Convolution Neural Network




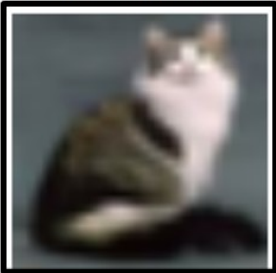




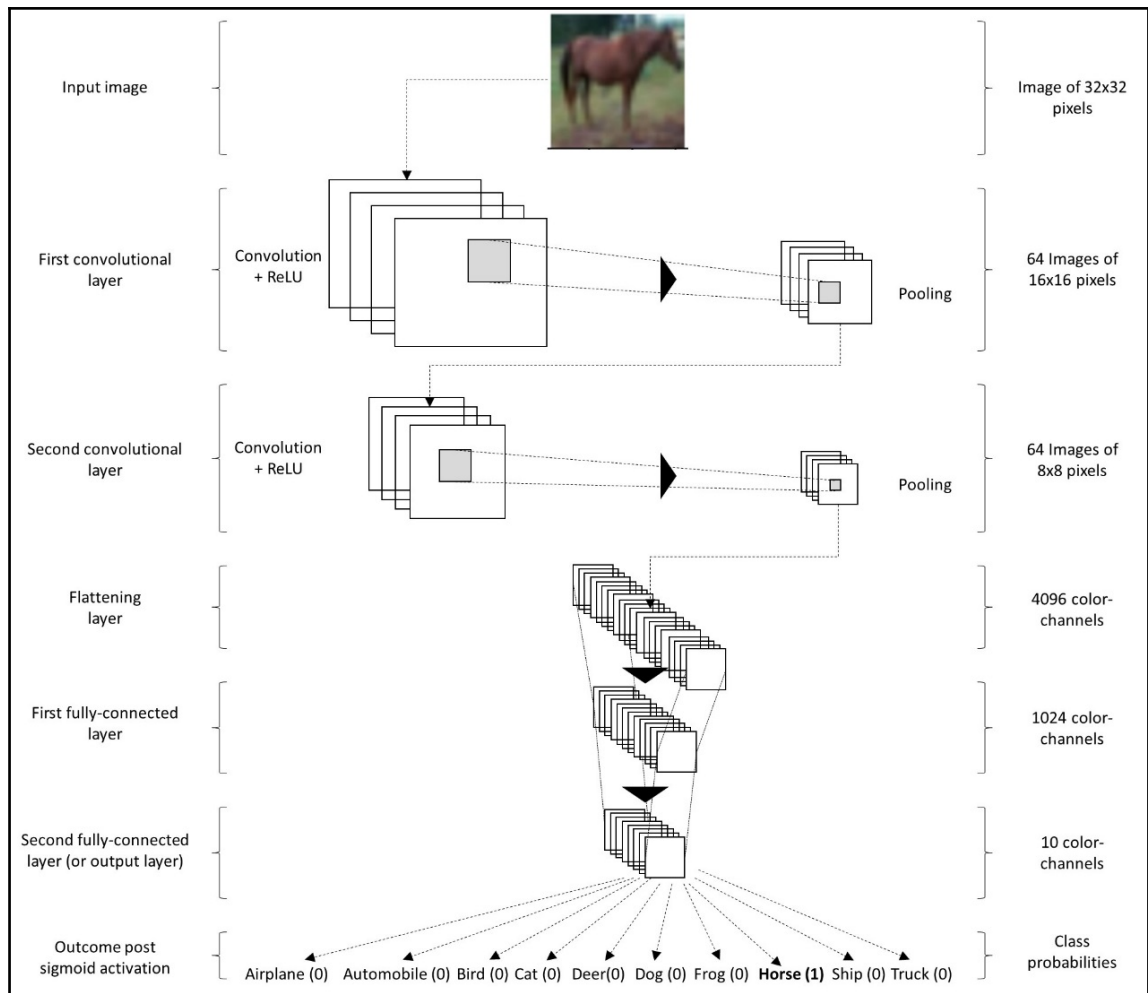




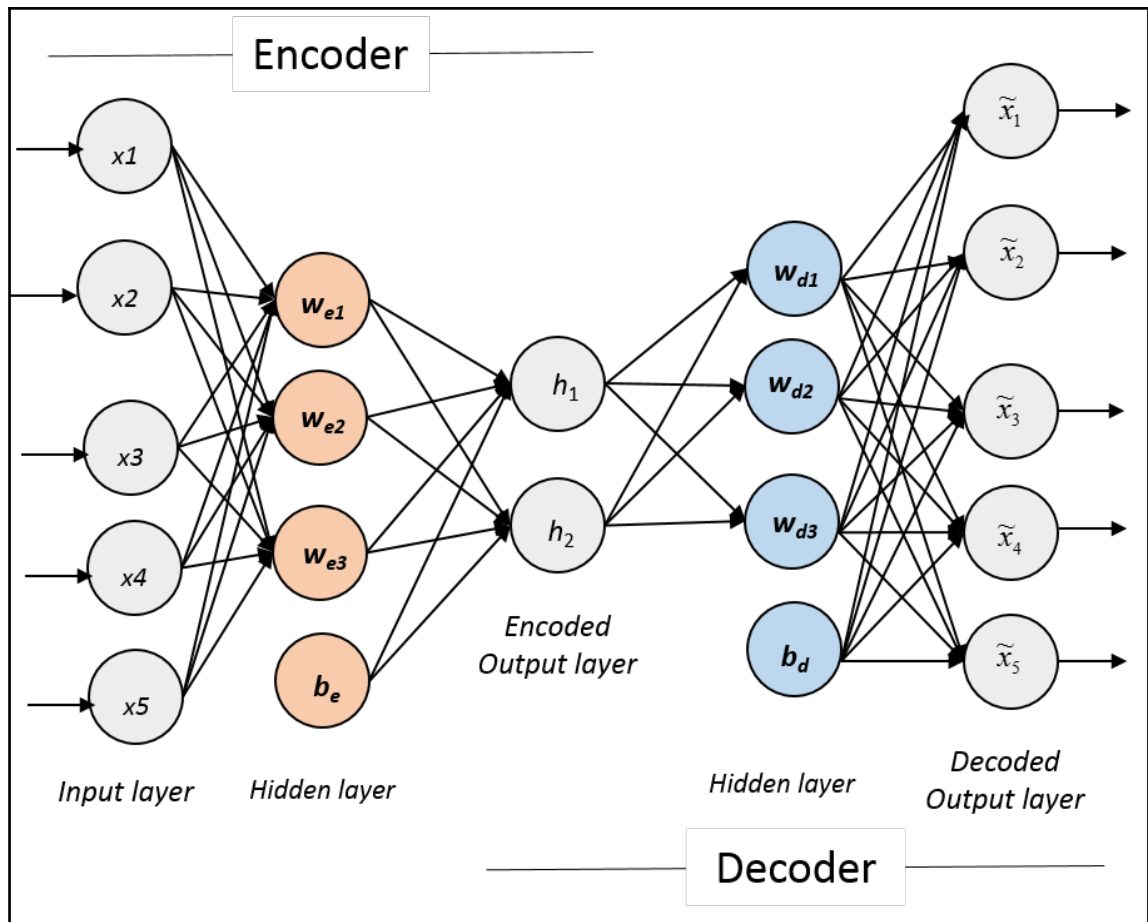
predicted										
actual	1	2	3	4	5	6	7	8	9	10
1	582	36	50	28	55	22	22	29	117	59
2	46	598	9	30	24	24	13	24	38	194
3	62	18	348	100	177	96	89	63	20	27
4	21	25	80	368	88	211	83	66	21	37
5	29	17	106	80	439	74	119	103	19	14
6	17	17	80	211	70	416	65	89	12	23
7	5	23	66	87	89	64	594	41	6	25
8	23	32	49	89	106	82	41	520	11	47
9	108	67	14	29	37	17	11	19	632	66
10	41	151	12	48	16	27	24	56	63	562

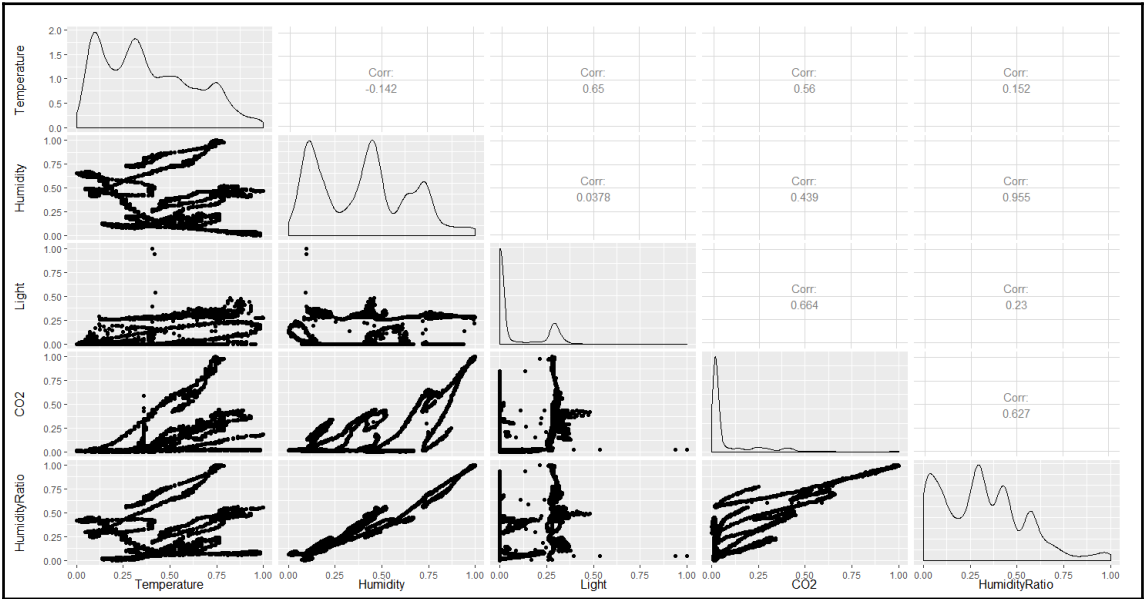
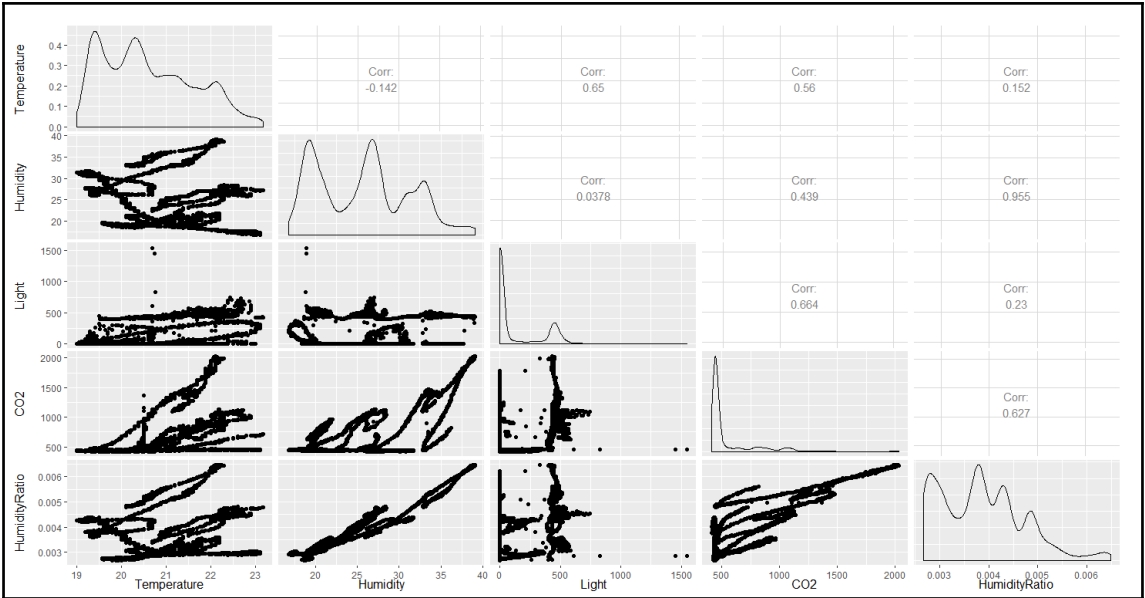


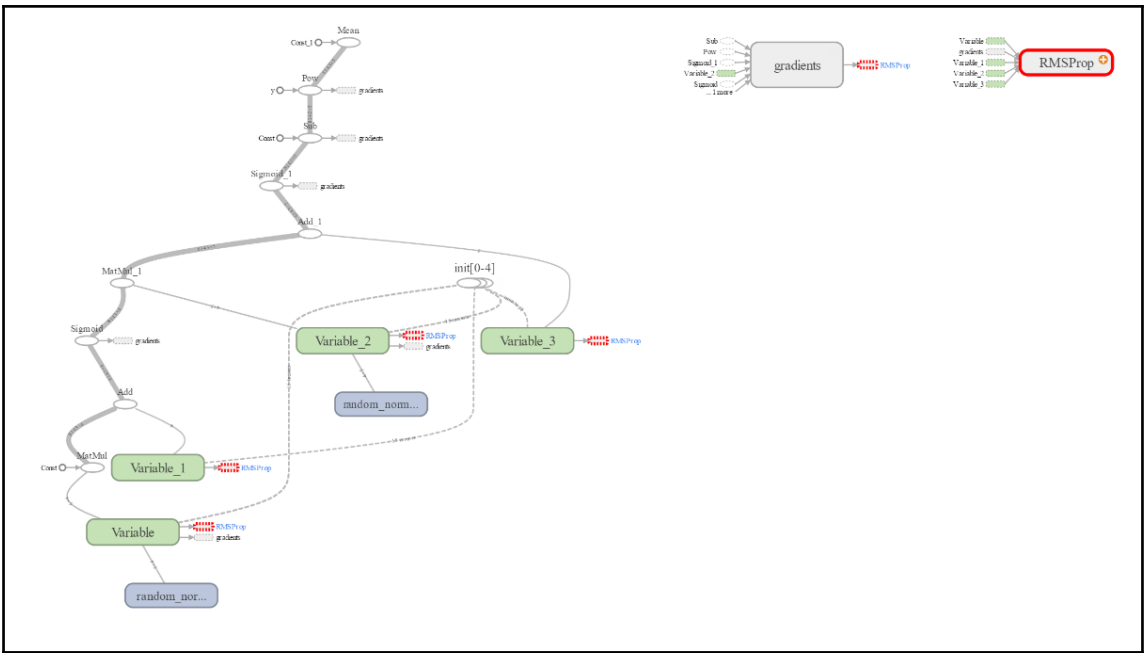
<p>True: automobile:: Pred: airplane</p> 	<p>True: horse:: Pred: deer</p> 	<p>True: airplane:: Pred: dog</p> 
<p>True: cat:: Pred: frog</p> 	<p>True: horse:: Pred: cat</p> 	<p>True: dog:: Pred: bird</p> 




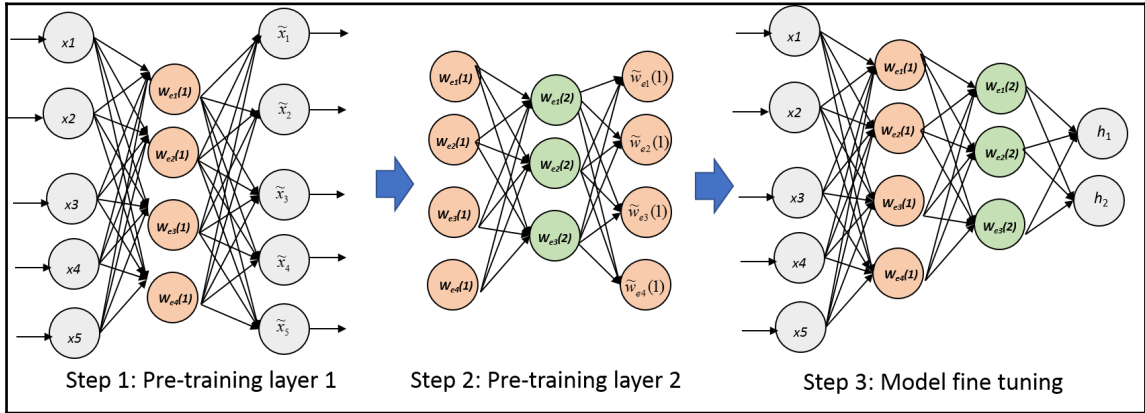
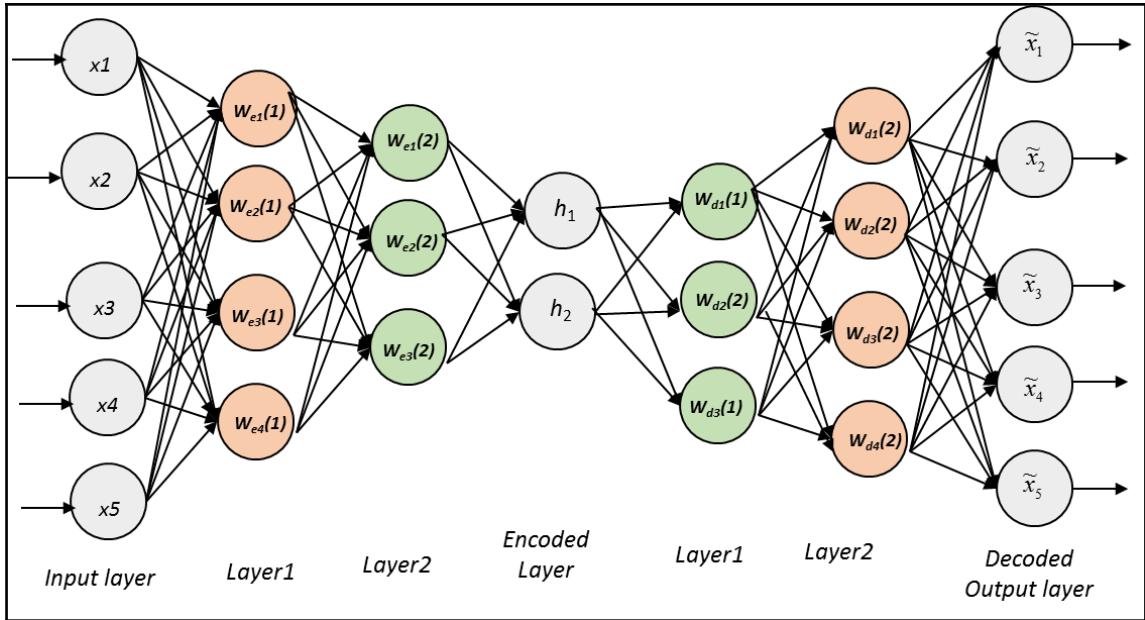
Chapter 4: Data Representation Using Autoencoders

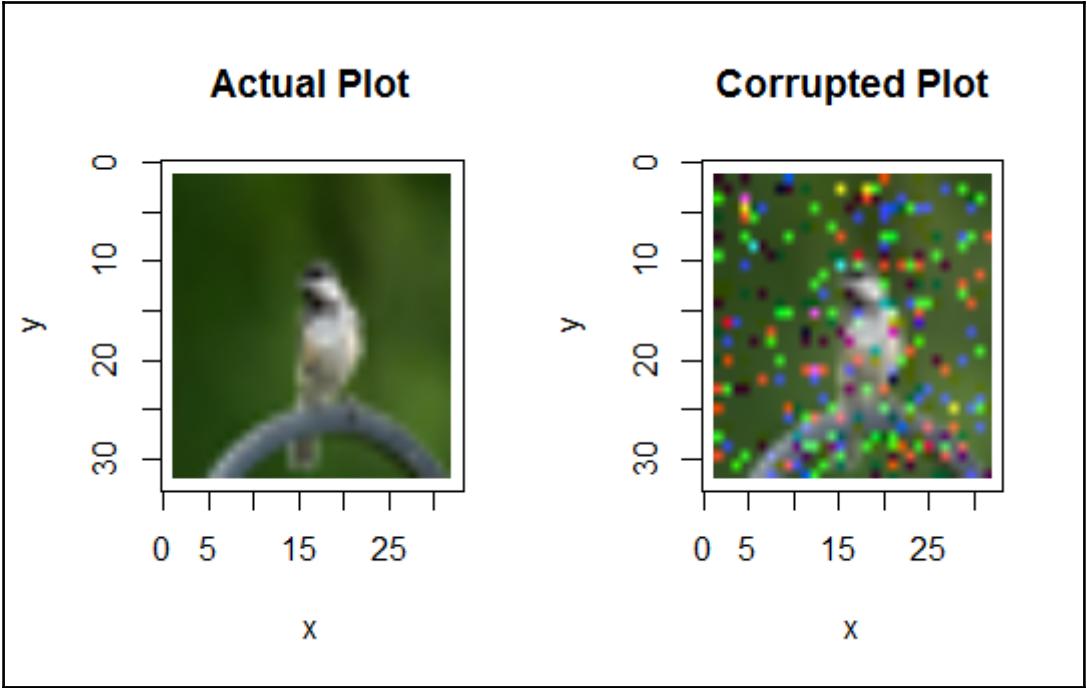
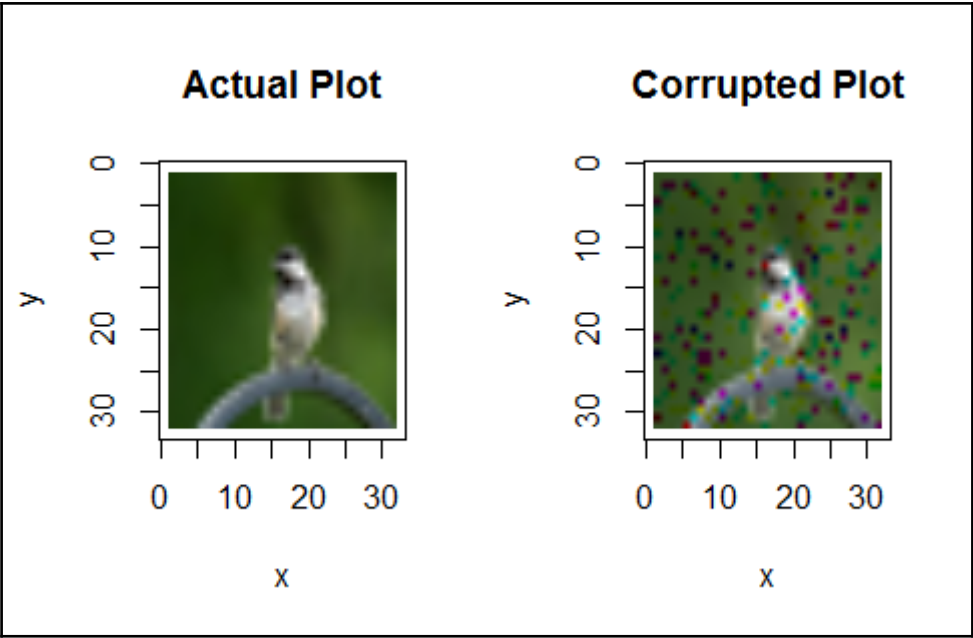


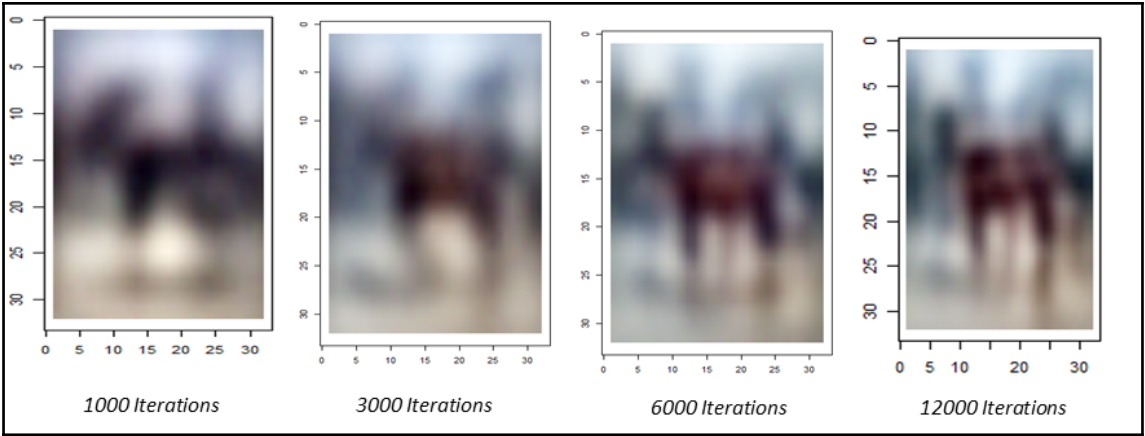



















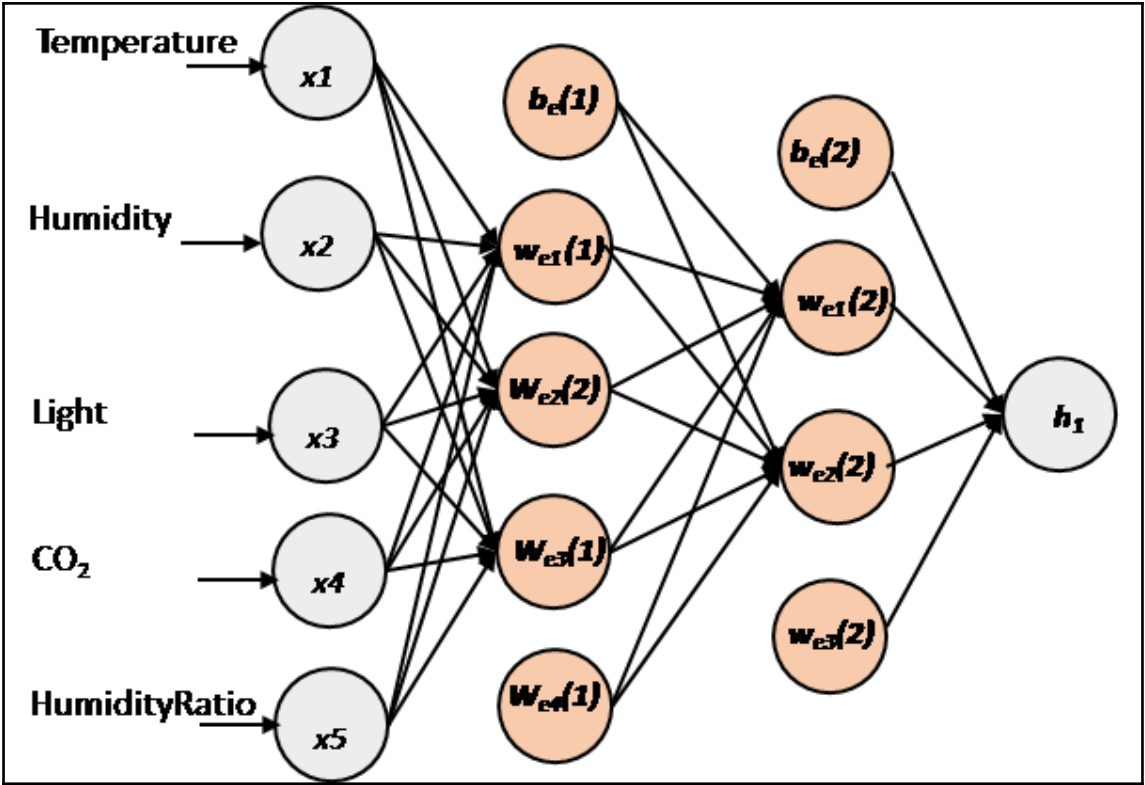



























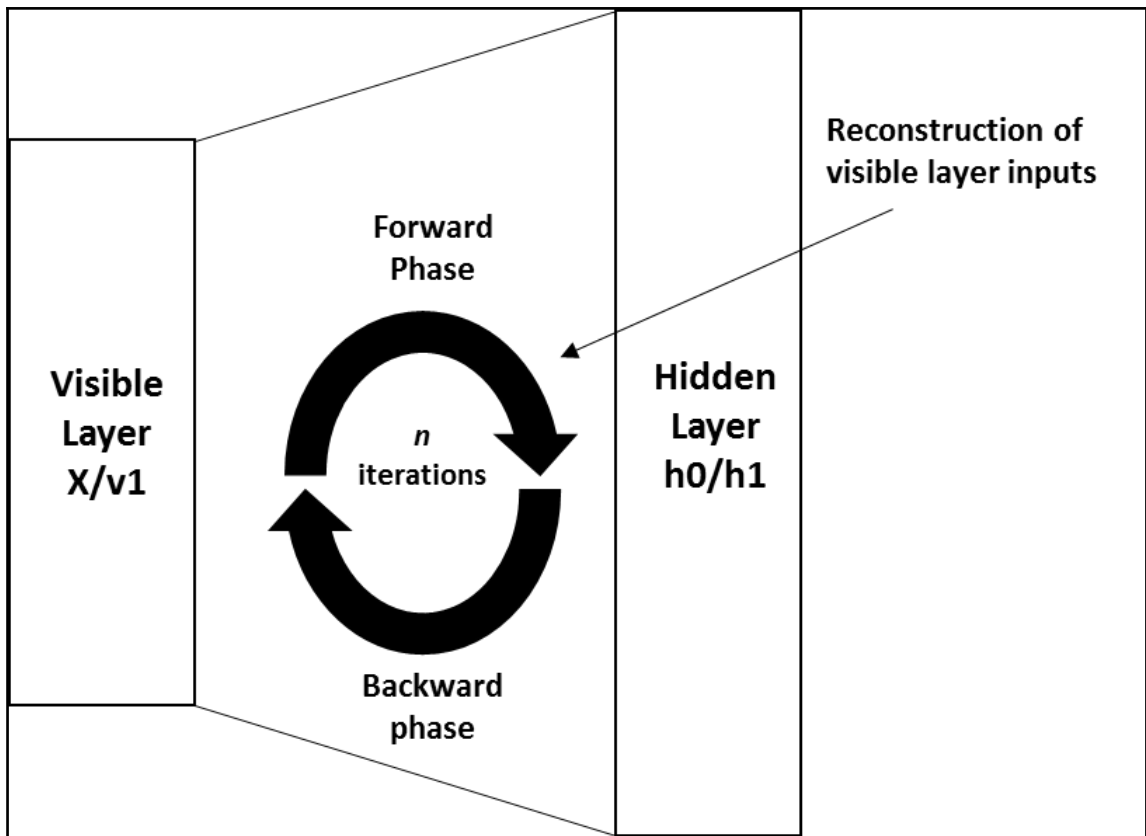
Chapter 5: Generative Models in Deep Learning

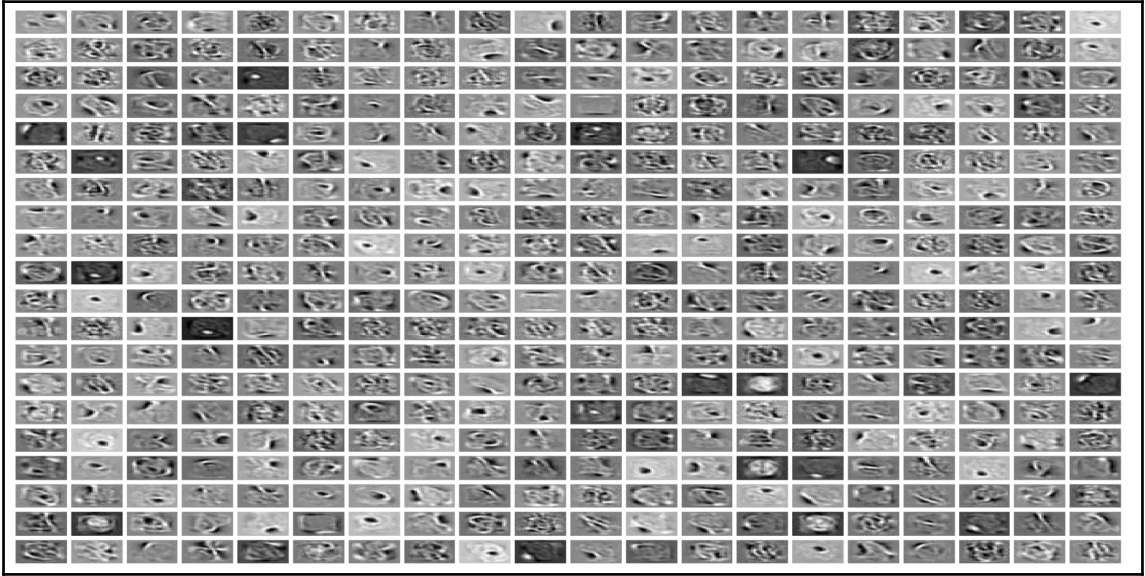






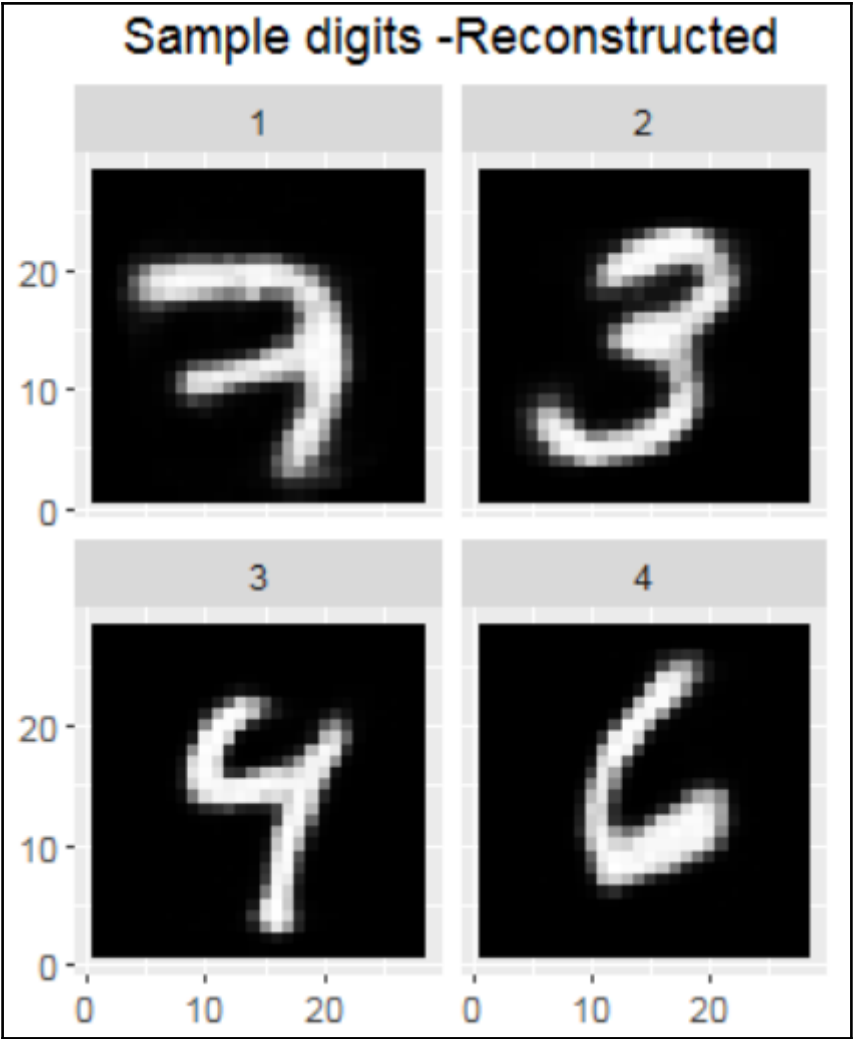




















```
> top_recom_movies
[1] "Star Wars: Episode VI - Return of the Jedi (1983)"
[2] "Matrix, The (1999)"
[3] "Star Wars: Episode V - The Empire Strikes Back (1980)"
[4] "Jurassic Park (1993)"
[5] "Star Wars: Episode IV - A New Hope (1977)"
[6] "Terminator 2: Judgment Day (1991)"
[7] "Raiders of the Lost Ark (1981)"
[8] "Star Wars: Episode I - The Phantom Menace (1999)"
[9] "Men in Black (1997)"
[10] "Princess Bride, The (1987)"
```


Type of Parameter	Parameter name	Parameter description
Input (Pre RBM)	<i>input_data</i>	Matrix of train MNIST data
Input (Pre RBM)	<i>num_input</i>	Number of independent variables
Input (Pre RBM)	<i>num_output</i>	Number of nodes in the corresponding hidden layer
Input (Pre RBM)	<i>epochs</i>	Number of iterations
Input (Pre RBM)	<i>alpha</i>	Learning rate for updating weight matrix
Input (Pre RBM)	<i>batchsize</i>	Number of observations per batch run
Output (Post RBM)	<i>output_data</i>	Matrix of reconstructed output
Output (Post RBM)	<i>error_list</i>	List of reconstruction errors for every run of 10 batch
Output (Post RBM)	<i>weight_list</i>	List of weight matrices for every run of 10 batch
Output (Post RBM)	<i>weight_final</i>	Matrix of final weights obtained after all epochs
Output (Post RBM)	<i>bias_final</i>	Vector of final biases obtained after all epochs



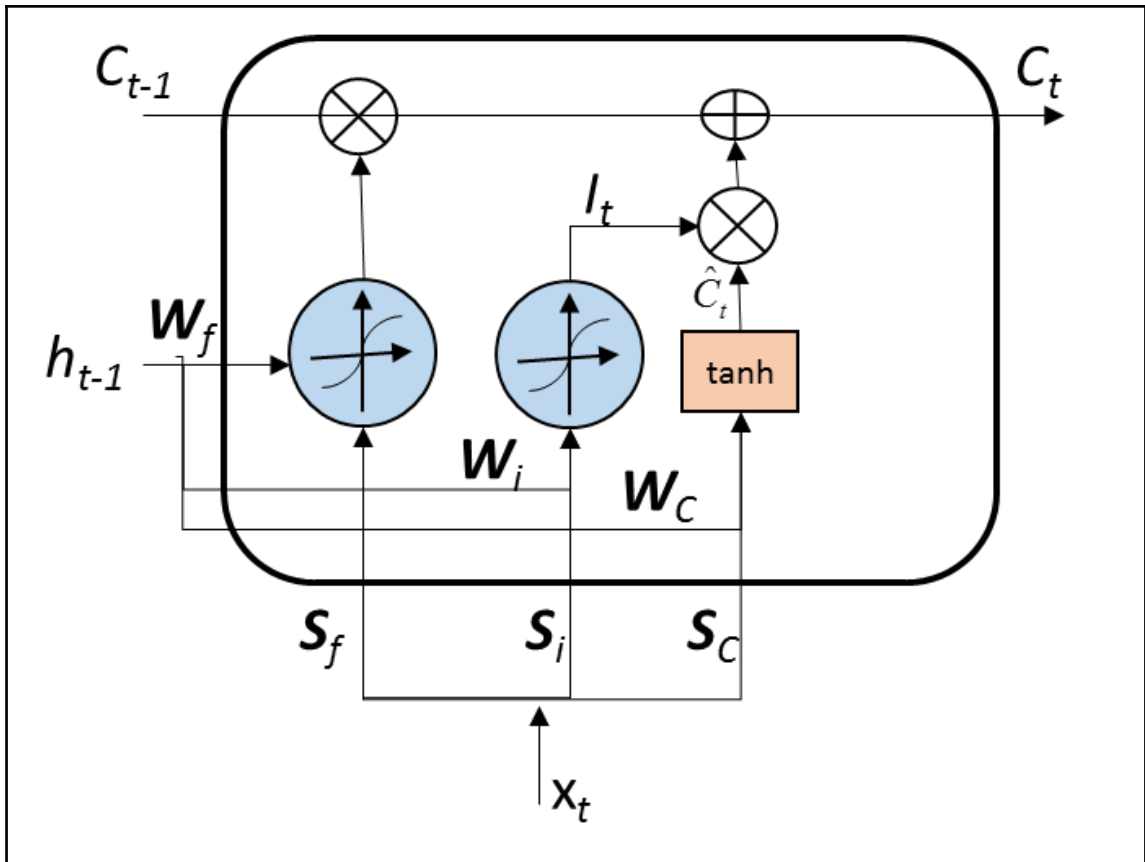
Parameter name	Parameter description
<i>Xdata</i>	Matrix of train input MNIST data
<i>Ydata</i>	Matrix of train output MNIST data
<i>Xtestdata</i>	Matrix of test input MNIST data
<i>Ytestdata</i>	Matrix of test output MNIST data
<i>input_size</i>	Number of attributes (or pixels) in train data
<i>learning_rate</i>	Learning rate for updating weight matrix
<i>momentum</i>	Increment in size of steps to jump out of local minima
<i>epochs</i>	Number of iterations
<i>batchsize</i>	Number of observations per batch run
<i>rbm_list</i>	List of outcomes of stacked RBM
<i>dbn_sizes</i>	Vector of sizes of hidden layers in stacked RBM


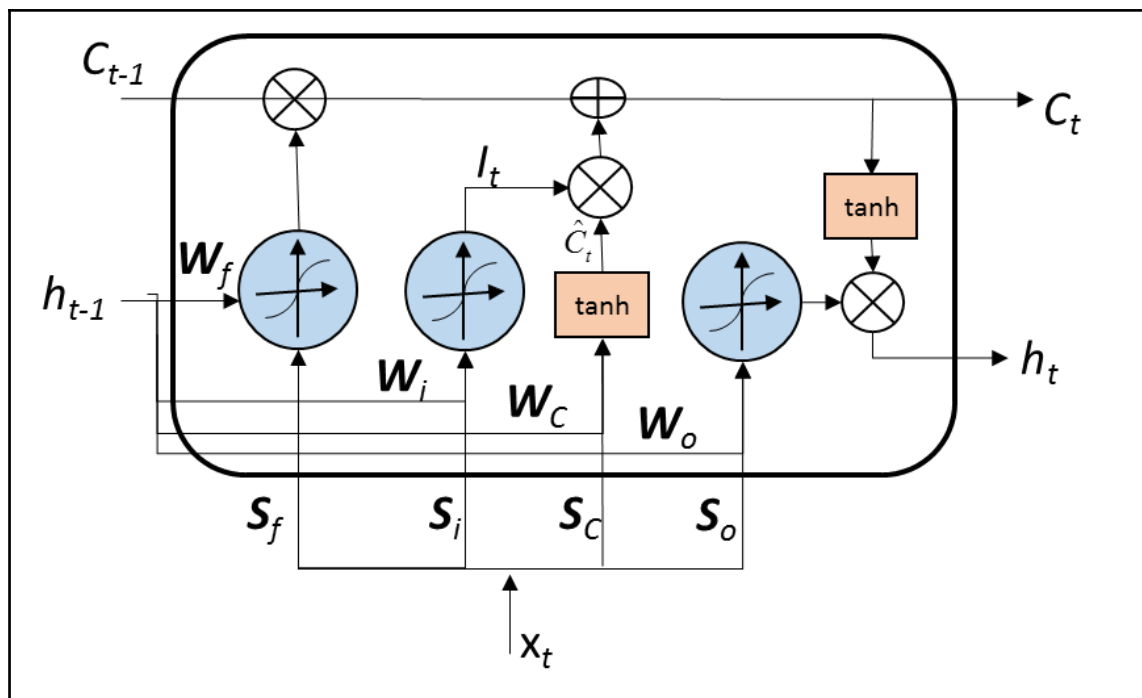


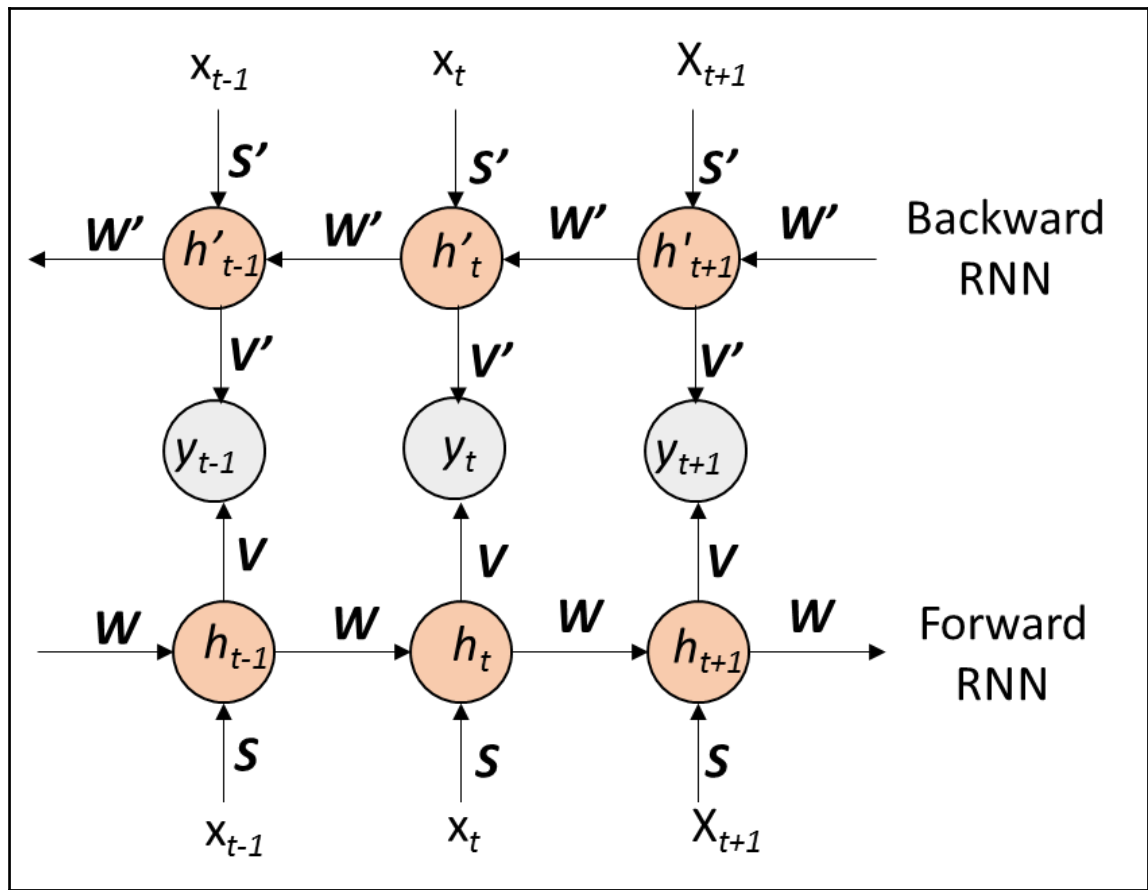
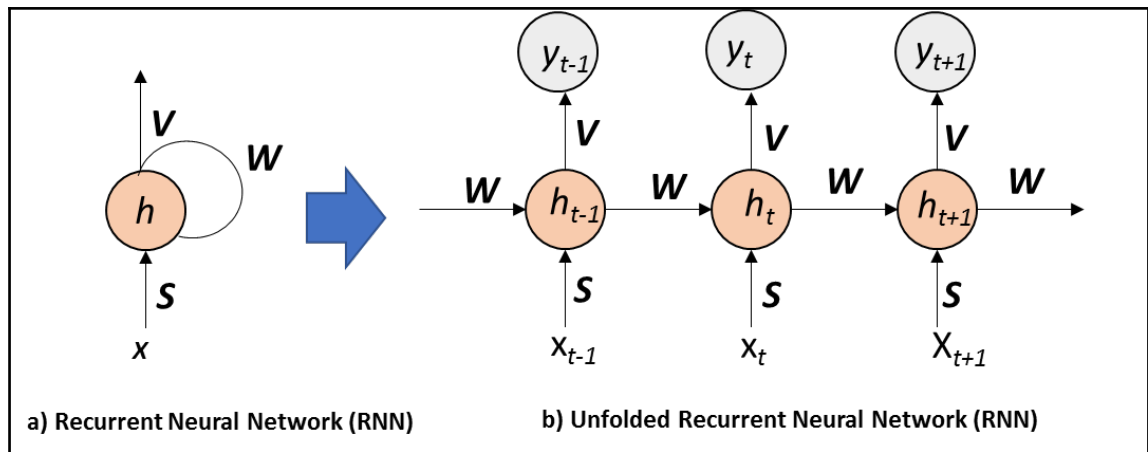





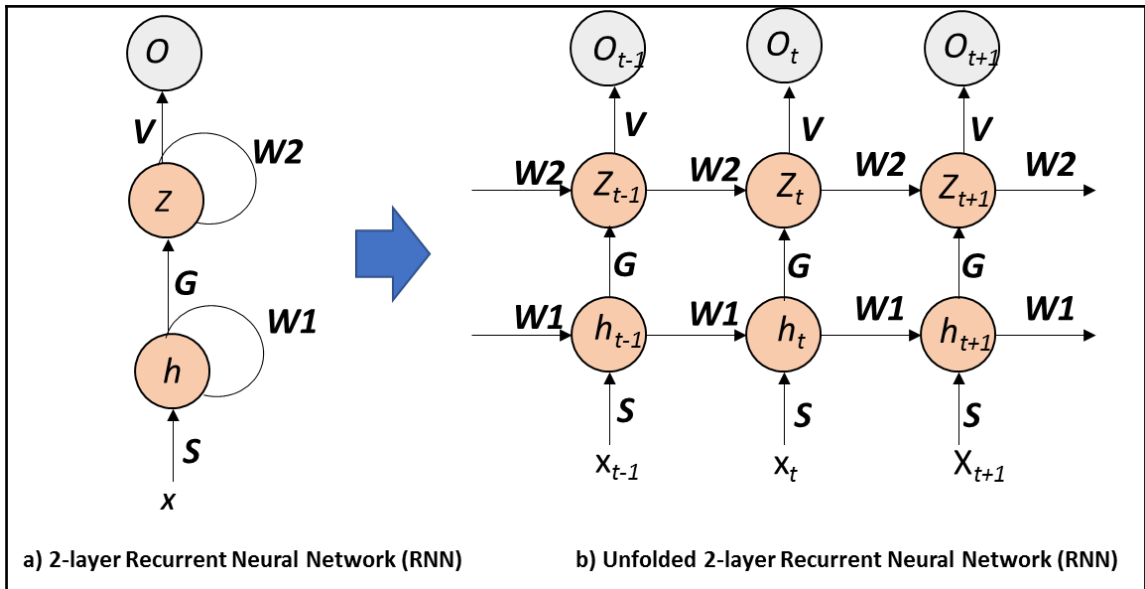
Chapter 6: Recurrent Neural Networks



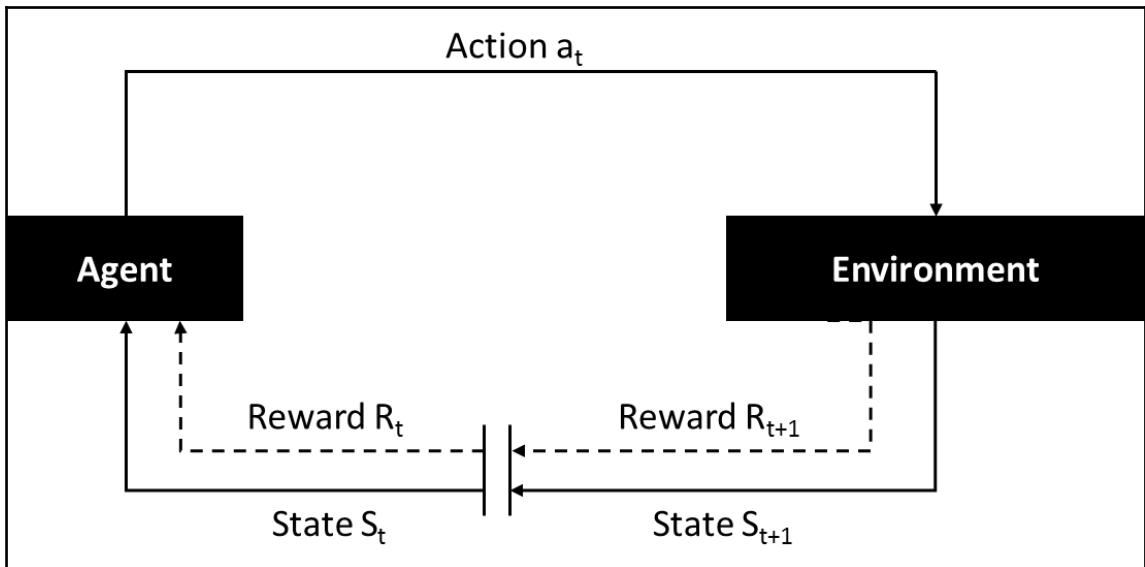




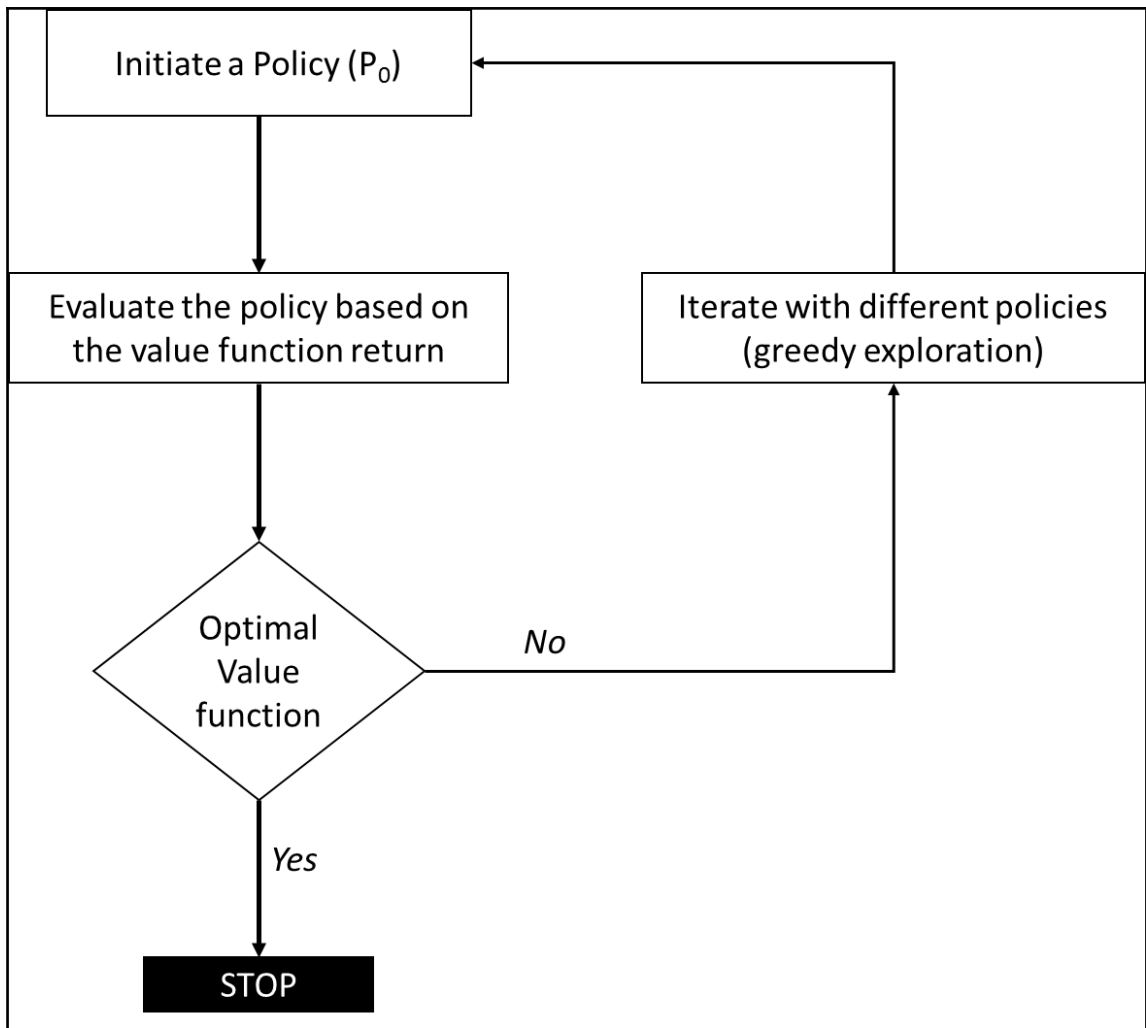





















Chapter 7: Reinforcement Learning

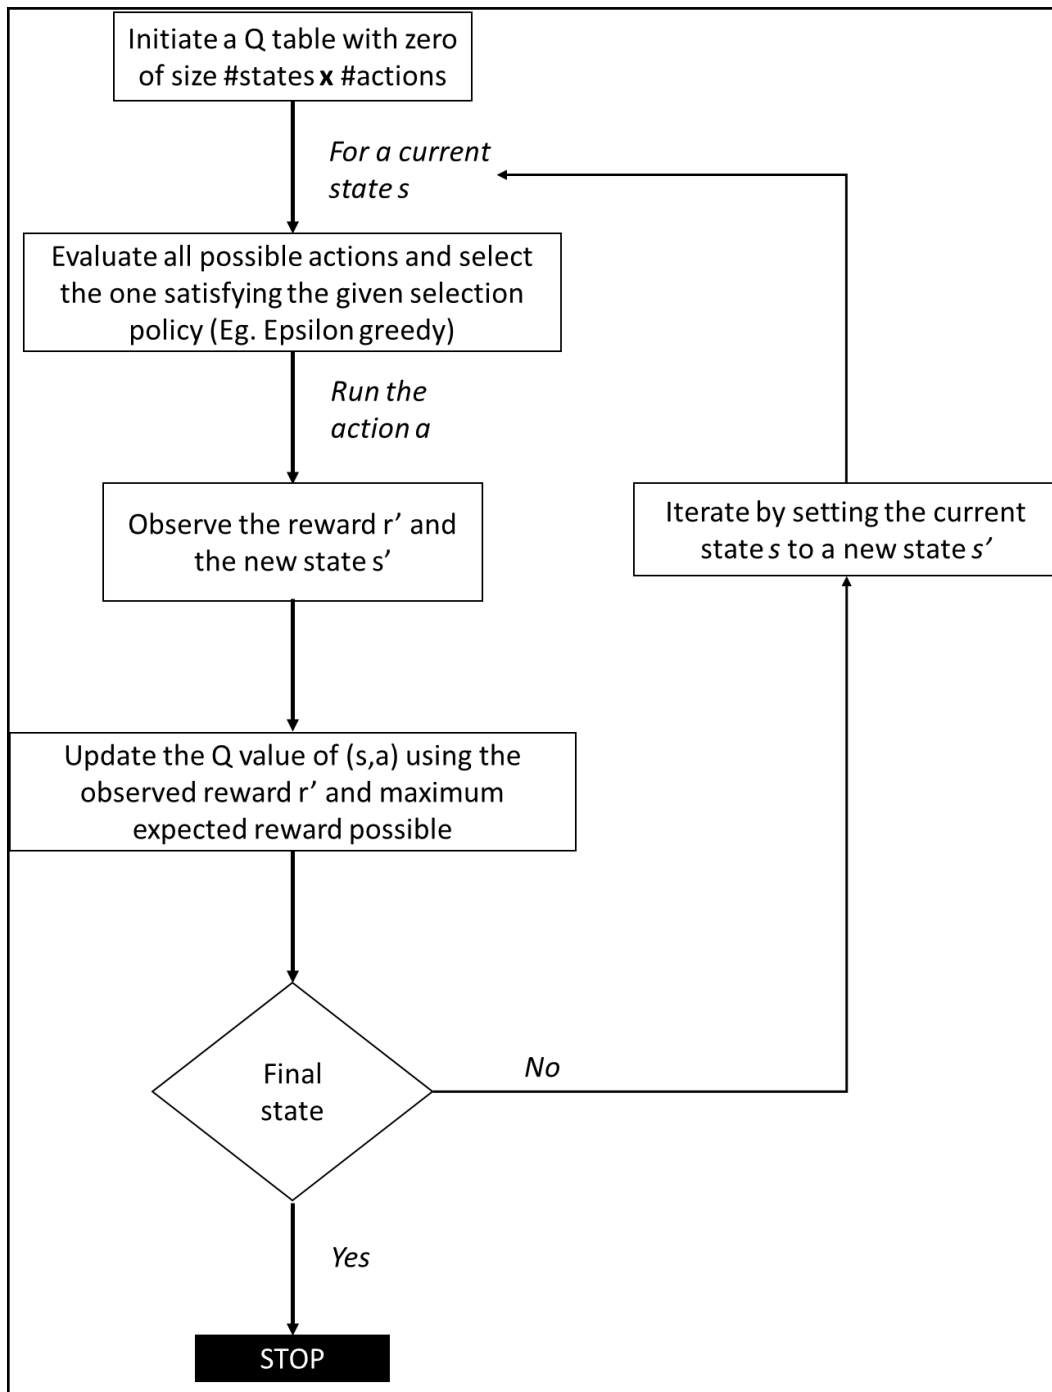


















S1 -1	S5 -1	S9 -1	S13 -1
S2 -1	S6 -1	S10 -1	S14 -1
S3 -1	S7 -1	S11 -1	S15 100
S4 -1	S8 -1	S12 -1	S16 -1





S1  -1	S5  -1	S9  -1	S13  -1
S2  -1	S6  -1	S10  -1	S14  -1
S3  -1	S7  -1	S11  -1	S15  100
S4  -1	S8  -1	S12  -1	S16  -1










S1  -1	S5  -1	S9  -1	S13  -1
S2  -1	S6  -1	S10  -1	S14  -1
S3  -1	S7  -1	S11  -1	S15  100
S4  -1	S8  -1	S12  -1	S16  -1

Chapter 8: Application of Deep Learning in Text Mining









Chapter 10: Transfer Learning

