


Chapter 1: Getting Started with Machine Learning for AWS



Data Gathering

- Obtain data from various sources
- Define schema to store data

Evaluation Metrics

- Define objectives of machine learning model
- Define evaluation metrics that determine success

Algorithm Selection

- Feature preprocessing
- Select algorithms to test
- Evaluate models on holdout datasets

Chapter 2: Classifying Twitter Feeds with Naive Bayes

The screenshot shows the Amazon SageMaker Dashboard. The left sidebar contains a navigation menu with categories: Dashboard, Search^{Beta}, Ground Truth (Labeling jobs, Labeling datasets, Labeling workforces), Notebook (Notebook instances, Lifecycle configurations, Git repositories), Training (Algorithms, Training jobs, Hyperparameter tuning jobs), and Inference (Compilation jobs, Model packages, Models, Endpoint configurations, Endpoints, Batch transform jobs). The main content area is titled 'Amazon SageMaker > Dashboard' and features an 'Overview' section with four cards: Ground Truth, Notebook, Training, and Inference. Each card has an icon, a brief description, and a button to access related resources. Below the Overview is a 'Recent activity' section with a filter for 'Last 7 days' and a message stating 'No recent activity.'

Git repositories Delete Update Secret Add repository

Search git repositories < 1 > ⚙️

Name	URL	ARN	Creation time
<input type="radio"/> mastering-ml-on-aws	https://github.com/mg-um/mastering-ml-on-aws.git	arn:aws:sagemaker:us-east-1:095585830284:code-repository/mastering-ml-on-aws	Dec 05, 2018 01:59 UTC

Notebook instance settings

Notebook instance name

Maximum of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an AWS Region.

Notebook instance type

Elastic Inference [Learn more](#)

IAM role

Notebook instances require permissions to call other services including SageMaker and S3. Choose a role or let us create a role with the `AmazonSageMakerFullAccess` IAM policy attached.

VPC - optional

Your notebook instance will be provided with SageMaker provided internet access because a VPC setting is not specified.

Lifecycle configuration - optional

Customize your notebook environment with default scripts and plugins.

Encryption key - optional

Encrypt your notebook data. Choose an existing KMS key or enter a key's ARN.

Volume Size In GB - optional

Your notebook instance's volume size in GB. Minimum of 5GB. Maximum of 16384GB (16TB).

▼ Git repositories - optional

▼ Default repository

Repository

Jupyter will start in this repository. Repositories are added to your home directory.





Files

Running

Clusters

SageMaker Examples

Conda

Select items to perform actions on them.

0



/ **mastering-ml-on-aws** / **chapter2**

..


collect_tweets.ipynb

train_scikit.ipynb

train_spark.ipynb


stop_words.txt

Chapter 3: Predicting House Value with Regression Algorithms



	ID	crim	zn	indus	chas	nox	rm	age	dis	rad	tax	ptratio	black	lstat	medv
0	1	0.00632	18.0	2.31	0	0.538	6.575	65.2	4.0900	1	296	15.3	396.90	4.98	24.0
1	2	0.02731	0.0	7.07	0	0.469	6.421	78.9	4.9671	2	242	17.8	396.90	9.14	21.6
2	4	0.03237	0.0	2.18	0	0.458	6.998	45.8	6.0622	3	222	18.7	394.63	2.94	33.4
3	5	0.06905	0.0	2.18	0	0.458	7.147	54.2	6.0622	3	222	18.7	396.90	5.33	36.2

	ID	crim	zn	indus	chas	nox	rm	age	dis	rad	tax	ptratio	black	lstat	medv	predicted_medv
74	110	0.26363	0.0	8.56	0	0.520	6.229	91.2	2.5451	5	384	20.9	391.23	15.55	19.4	20.058458
310	473	3.56868	0.0	18.10	0	0.580	6.437	75.0	2.8965	24	666	20.2	393.37	14.36	23.2	20.850997
264	404	24.80170	0.0	18.10	0	0.693	5.349	96.0	1.7028	24	666	20.2	396.90	19.77	8.3	13.150412
207	311	2.63548	0.0	9.90	0	0.544	4.973	37.8	2.5194	4	304	18.4	350.45	12.64	16.1	18.640210
283	440	9.39063	0.0	18.10	0	0.740	5.627	93.9	1.8172	24	666	20.2	396.90	22.88	12.8	11.553577



ID	crim	zn	indus	chas	nox	rm	age	dis	rad	tax	prratio	black	lstat	medv
1	0.00632	18.0	2.31	0	0.538	6.575	65.2	4.09	1	296	15.3	396.9	4.98	24.0
2	0.02731	0.0	7.07	0	0.469	6.421	78.9	4.9671	2	242	17.8	396.9	9.14	21.6
4	0.03237	0.0	2.18	0	0.458	6.998	45.8	6.0622	3	222	18.7	394.63	2.94	33.4
5	0.06905	0.0	2.18	0	0.458	7.147	54.2	6.0622	3	222	18.7	396.9	5.33	36.2


ID	crim	zn	indus	chas	nox	rm	age	dis	rad	tax	prratio	black	lstat	medv	features
1	0.00632	18.0	2.31	0	0.538	6.575	65.2	4.09	1	296	15.3	396.9	4.98	24.0	[0.00632,18.0,2.3...
2	0.02731	0.0	7.07	0	0.469	6.421	78.9	4.9671	2	242	17.8	396.9	9.14	21.6	[0.02731,0.0,7.07...
4	0.03237	0.0	2.18	0	0.458	6.998	45.8	6.0622	3	222	18.7	394.63	2.94	33.4	[0.03237,0.0,2.18...

ID	crim	zn	indus	chas	nox	rm	age	dis	rad	tax	ptratio	black	lstat	medv	features	prediction
7	0.08829	12.5	7.87	0	0.524	6.012	66.6	5.5605	5	311	15.2	395.6	12.43	22.9	[0.08829,12.5,7.8...	21.273530243958177
24	0.98843	0.0	8.14	0	0.538	5.813	100.0	4.0952	4	307	21.0	394.54	19.88	14.5	[0.98843,0.0,8.14...	13.894245541490553
44	0.15936	0.0	6.91	0	0.448	6.211	6.5	5.7209	3	233	17.9	394.46	7.44	24.7	[0.15936,0.0,6.91...	25.209683694484067

score	
0	18.911674
1	41.916256
2	20.833599
3	38.696209
4	30.833647

	actual	predicted
0	14.3	18.911674
1	50.0	41.916256
2	23.2	20.833599
3	46.0	38.696209
4	30.8	30.833647

Chapter 4: Predicting User Behavior with Tree-Based Methods



Name your notebook, choose a cluster or create one, and customize configuration options if desired. [Learn more](#)

Notebook name*

Names may only contain letters (a-z), numbers (0-9), hyphens (-), or underscores (_).

Description

256 characters max.

Cluster* Choose an existing cluster

Create a cluster ?

Cluster name:

Release: emr-5.20.0

Applications: Hadoop, Spark, Livy

Instance:

EMR role: [EMR_DefaultRole](#) ?

EC2 instance profile: [EMR_EC2_DefaultRole](#) ?

Security groups Use default security groups ?

Choose security groups (vpc-8e36a0f5)

AWS service role* ?

Notebook location* Choose an S3 location in us-east-1 ?

▶ **Tags** ?

* Required

Cancel

Create notebook

Notebook: emr_notebook Pending Notebook is attaching to cluster j-3CL4E6JJWJ7GJ. Notebook can now be used in local mode.


[Open](#) [Stop](#) [Delete](#)

Notebook C












Notebook ID: e-2SNYVRZPK67OMAZZCPESUOXWJ
Description: --
Last modified: 1 second ago ⓘ
Last modified by: ...root ⓘ
Created on: 2019-05-08 18:25 (UTC-4)
Created by: ...root ⓘ
Service IAM role: [EMR_Notebooks_DefaultRole](#) ⓘ
Notebook tags: creatorUserId = 095585830284 [View All / Edit](#)
Notebook location: s3://mastering-ml-aws/sample-notebook/ ⓘ

Cluster

Cluster: NotebookCluster
Cluster Id: [j-3CL4E6JJWJ7GJ](#)
Cluster status: Starting
Cluster tags: creator = NOTEBOOK_CONSOLE [View All](#)
Step logs: s3://aws-logs-095585830284-us-east-1/elasticmapreduce/ ⓘ

 **jupyter** emr_notebook (unsaved changes)

File Edit View Insert Cell Kernel Widgets Help Trusted ⓘ | Kernel O

           Code ⓘ

```
In [ ]: |
```

aws Services Resource Groups montevidelabs N. Virginia Support

Create Cluster - Advanced Options [Go to quick options](#)

Step 1: Software and Steps

Step 2: Hardware

Step 3: General Cluster Settings

Step 4: Security

Software Configuration

Release: emr-5.23.0

<input checked="" type="checkbox"/> Hadoop 2.8.5	<input type="checkbox"/> Zeppelin 0.8.1	<input type="checkbox"/> Livy 0.5.0
<input type="checkbox"/> JupyterHub 0.9.4	<input type="checkbox"/> Tez 0.9.1	<input type="checkbox"/> Flink 1.7.1
<input checked="" type="checkbox"/> Ganglia 3.7.2	<input type="checkbox"/> HBase 1.4.9	<input type="checkbox"/> Pig 0.17.0
<input checked="" type="checkbox"/> Hive 2.3.4	<input type="checkbox"/> Presto 0.215	<input type="checkbox"/> ZooKeeper 3.4.13
<input type="checkbox"/> MXNet 1.3.1	<input type="checkbox"/> Sqoop 1.4.7	<input type="checkbox"/> Mahout 0.13.0
<input type="checkbox"/> Hue 4.3.0	<input type="checkbox"/> Phoenix 4.14.1	<input type="checkbox"/> Oozie 5.1.0
<input checked="" type="checkbox"/> Spark 2.4.0	<input type="checkbox"/> HCatalog 2.3.4	<input type="checkbox"/> TensorFlow 1.12.0

AWS Glue Data Catalog settings (optional)

Use for Hive table metadata

Use for Spark table metadata

Edit software settings

Enter configuration Load JSON from S3

```
[[{"classification": "livy-conf", "properties": {"livy.server.session.timeout": "12h"}}]]
```

```
In [1]: s3_train_path = 's3://mastering-ml-aws/chapter4/training-data'
```

VBox()

Starting Spark application

ID	YARN Application ID	Kind	State	Spark UI	Driver log	Current session?
1	application_1556298146643_0008	pyspark	idle	Link	Link	✓

SparkSession available as 'spark'.

c0	c1	c2	c3	c4	c5	c6	c7	c8	c9	l
[-1664374510]	[1292560685]	[1963151207]	[-113426919]	[1024827180]	null	[-1841755489, -20...	[781804810]	[677061876]	[-2054476127, 128...	0
[1566608579]	[-248982458]	[336746857]	[-1629610286]	[244157766]	null	[-574085389, 1869...	[1065163157]	[332083152]	[-614983515]	0
[1935105702]	[1292560685]	[-1389162932]	[-113426919]	[-8361123]	[839761088]	null	[-1708330775]	[1856995055]	[-1954958362, 157...	0
[1718276659]	[630920017]	[1171414431]	[-113426919]	[640993460]	[-1183679474]	[-833071846, 1997...	[680566046]	[-681791195]	[-813776566, 1925...	1
[1562430026]	[630920017]	[1639152385]	[1781226914]	[1493440023]	null	[-2043098156, -95...	[324520841]	[873551722]	[980231370]	0

f0	f1	f2	f3	f4	click
-1664374510	1292560685	1963151207	-113426919	1024827180	0
1566608579	-248982458	336746857	-1629610286	244157766	0
1935105702	1292560685	-1389162932	-113426919	-8361123	0
1718276659	630920017	1171414431	-113426919	640993460	1
1562430026	630920017	1639152385	1781226914	1493440023	0

summary	f0	f1	f2	f3	f4	click
count	12000000	12000000	12000000	12000000	12000000	12000000
mean	-6.610412663970825E7	2.5049429668800482E8	-2.915904354482062E8	5.459869260236725E7	-6.716129061083934E7	0.18310175
stddev	1.2294656059145813E9	1.287445524252859E9	1.2580392622053525E9	8.234483651283175E8	1.2429134469135067E9	0.38674993421016063
min	-2145952914	-2125813709	-2145112401	-2134594413	-2147400218	0
max	2146734164	2136145316	2145529900	2102865870	2147086554	1

```
df.select("f0").distinct().count()
```

```
2497
```

```
df.select("f1").distinct().count()
```

```
178
```

```
df.select("f3").distinct().count()
```

```
68
```

```
df.select("f4").distinct().count()
```

```
17572
```

```
+-----+-----+
|          f0 | f0_index |
+-----+-----+
| -130745722 |    178.0 |
| -1322326169 |     0.0 |
| -130745722 |    178.0 |
| -1248885727 |     8.0 |
|  571589560 |    877.0 |
+-----+-----+
only showing top 5 rows
```



```

+-----+-----+-----+
|          f0 | f0_index |          f0_encoded |
+-----+-----+-----+
| -1910840705 |    118.0 | (2496,[118],[1.0]) |
| -1713169383 |    242.0 | (2496,[242],[1.0]) |
|  1590237751 |   1216.0 | (2496,[1216],[1.0]) |
| -1156005499 |    337.0 | (2496,[337],[1.0]) |
|  1707433888 |    388.0 | (2496,[388],[1.0]) |
+-----+-----+-----+

```

only showing top 5 rows

2.4.0

Jobs
Stages
Storage
Environment
Executors
SQL

livy-session-0 application UI


Spark Jobs (?)


User: livy
Total Uptime: 7.3 min
Scheduling Mode: FIFO
Active Jobs: 1
Completed Jobs: 2

▶ Event Timeline

▼ **Active Jobs (1)**

Job Id (Job Group)	Description	Submitted	Duration	Stages: Succeeded/Total	Tasks (for all stages): Succeeded/Total
2 (6)	countByValue at StringIndexer.scala:140 countByValue at StringIndexer.scala:140 (kill)	2019/01/26 14:31:10	42 s	0/2	0/2 (1 running)







selected_features	probability	rawPrediction
(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 1.0, ...	[0.8323081791819823, 0.16769182081801767]	[2677329.0, 539423.0]
(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...	[0.8323081791819823, 0.16769182081801767]	[2677329.0, 539423.0]
(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...	[0.8323081791819823, 0.16769182081801767]	[2677329.0, 539423.0]
(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...	[0.8323081791819823, 0.16769182081801767]	[2677329.0, 539423.0]

	click	f0	f1	f2	f3	f4	f5	f6	f7	f8	...	f91	f92	f93	f94	f95	f96	f97	f98	f99	score	
0	1	0.0	0.0	0.0	0.0	1.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.226555
1	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.167860
2	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.512123
3	1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.512123
4	1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.141251

Chapter 5: Customer Segmentation Using Clustering Algorithms






InvoiceNo	StockCode	Description	Quantity	InvoiceDate	UnitPrice	CustomerID	Country
0	536365 85123A	WHITE HANGING HEART T-LIGHT HOLDER	6	12/1/2010 8:26	2.55	17850.0	United Kingdom
1	536365 71053	WHITE METAL LANTERN	6	12/1/2010 8:26	3.39	17850.0	United Kingdom
2	536365 84406B	CREAM CUPID HEARTS COAT HANGER	8	12/1/2010 8:26	2.75	17850.0	United Kingdom
3	536365 84029G	KNITTED UNION FLAG HOT WATER BOTTLE	6	12/1/2010 8:26	3.39	17850.0	United Kingdom
4	536365 84029E	RED WOOLLY HOTTIE WHITE HEART.	6	12/1/2010 8:26	3.39	17850.0	United Kingdom

InvoiceNo	StockCode	Description	Quantity	InvoiceDate	UnitPrice	CustomerID	Country	TotalBought
0	536365	85123A	WHITE HANGING HEART T-LIGHT HOLDER	6	12/1/2010 8:26	2.55	17850 United Kingdom	15.30
1	536365	71053	WHITE METAL LANTERN	6	12/1/2010 8:26	3.39	17850 United Kingdom	20.34
2	536365	84406B	CREAM CUPID HEARTS COAT HANGER	8	12/1/2010 8:26	2.75	17850 United Kingdom	22.00
3	536365	84029G	KNITTED UNION FLAG HOT WATER BOTTLE	6	12/1/2010 8:26	3.39	17850 United Kingdom	20.34
4	536365	84029E	RED WOOLLY HOTTIE WHITE HEART.	6	12/1/2010 8:26	3.39	17850 United Kingdom	20.34


CustomerID	SumTotalBought	
0	17420	598.83
1	16861	151.65
2	16503	1421.43
3	15727	5178.96
4	17389	31300.08

CustomerID	InvoiceNo	StockCode	Description	Quantity	InvoiceDate	UnitPrice	Country	TotalBought	SumTotalBought	
0	17850	536365	85123A	WHITE HANGING HEART T-LIGHT HOLDER	6	12/1/2010 8:26	2.55	United Kingdom	15.30	5288.63
1	17850	536365	71053	WHITE METAL LANTERN	6	12/1/2010 8:26	3.39	United Kingdom	20.34	5288.63
2	17850	536365	84406B	CREAM CUPID HEARTS COAT HANGER	8	12/1/2010 8:26	2.75	United Kingdom	22.00	5288.63
3	17850	536365	84029G	KNITTED UNION FLAG HOT WATER BOTTLE	6	12/1/2010 8:26	3.39	United Kingdom	20.34	5288.63

TotalBought	SumTotalBought	stock_code_index	stock_code_encoded	country_index	country_encoded	total_bought_index	features_raw	features	prediction
15.30	5288.63	0.0	(1.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	0.0	(1.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	2.0	(1.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	(0.25, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	0
20.34	5288.63	403.0	(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	0.0	(1.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	2.0	(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	0
22.00	5288.63	452.0	(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	0.0	(1.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	2.0	(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	(0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, ...)	0



Mean total amount bought per cluster





_encoded	country_index	country_encoded	total_bought_index	features_raw	features	distance_to_cluster	closest_cluster
,[1.0]]	0.0	(37,[0],[1.0])	2.0	(3722,[0,3684,372...]	(3722,[0,3684,372...]	0.26693078875541687	2.0
,[1.0]]	0.0	(37,[0],[1.0])	2.0	(3722,[403,3684,3...]	(3722,[403,3684,3...]	0.2682397961616516	2.0
,[1.0]]	0.0	(37,[0],[1.0])	2.0	(3722,[452,3684,3...]	(3722,[452,3684,3...]	0.268250435590744	2.0
,[1.0]]	0.0	(37,[0],[1.0])	2.0	(3722,[288,3684,3...]	(3722,[288,3684,3...]	0.268185555934906	2.0
,[1.0]]	0.0	(37,[0],[1.0])	2.0	(3722,[281,3684,3...]	(3722,[281,3684,3...]	0.2680809795856476	2.0

Cluster: NotebookCluster Terminated Terminated by user request

Summary Application history Monitoring Hardware Events Steps Configurations Bootstrap actions

Connections: --

Master public DNS: ec2-54-224-42-189.compute-1.amazonaws.com SSH

Tags: creator = NOTEBOOK_CONSOLE View All

Summary	Configuration details	Network and hardware
ID: j-ORKC7X8NUP2M Creation date: 2019-02-14 11:18 (UTC-5) End date: 2019-02-14 12:51 (UTC-5) Elapsed time: 1 hour, 33 minutes Auto-terminate: No Termination protection: Off	Release label: emr-5.20.0 Hadoop distribution: Amazon 2.8.5 Applications: Spark 2.4.0, JupyterHub 0.9.4, Ganglia 3.7.2 Log URI: s3://aws-logs-095585830284-us-east-1/elasticmapreduce/ EMRFS consistent view: Disabled Custom AMI ID: --	Availability zone: us-east-1c Subnet ID: subnet-224d747f Master: Terminated 1 m5.xlarge Spot (max on-demand) Core: -- Task: --

Security and access

Key name: mlabskeys

EC2 instance profile: EMR_EC2_DefaultRole

EMR role: EMR_DefaultRole

Visible to all users: All Change

Security groups for sg-0f1c57cc06ee7773a Master: (ElasticMapReduce-master)

Security groups for sg-06063324b23c5501a Core & Task: (ElasticMapReduce-slave)

Role ARN `arn:aws:iam::095585830284:role/EMR_EC2_DefaultRole` [🔗](#)
Role description [Edit](#)
Instance Profile ARNs `arn:aws:iam::095585830284:instance-profile/EMR_EC2_DefaultRole` [🔗](#)
Path /
Creation time 2019-01-14 20:59 EST
Maximum CLI/API session duration 1 hour [Edit](#)

Permissions | **Trust relationships** | **Tags** | **Access Advisor** | **Revoke sessions**

▼ Permissions policies (3 policies applied)

[Attach policies](#) [+ Add inline policy](#)

Policy name ▼	Policy type ▼
▼ sagemaker-all	Managed policy ✕

[Policy summary](#) | [{} JSON](#) | [Edit policy](#) [Simulate policy](#)

```

1 {
2   "Version": "2012-10-17",
3   "Statement": [
4     {
5       "Sid": "VisualEditor0",
6       "Effect": "Allow",
7       "Action": "sagemaker:*",
8       "Resource": "*"
9     }
10  ]
11 }
  
```

Permissions | **Trust relationships** | **Tags** | **Access Advisor** | **Revoke sessions**

You can view the trusted entities that can assume the role and the access conditions for the role. [Show policy document](#)

[Edit trust relationship](#)

<p>Trusted entities</p> <p>The following trusted entities can assume this role.</p> <div style="border: 1px solid #ccc; background-color: #fff9c4; padding: 5px; margin: 5px 0;"> <p>Trusted entities</p> <p>The identity provider(s) <code>ec2.amazonaws.com</code></p> <p>The identity provider(s) <code>sagemaker.amazonaws.com</code></p> </div>	<p>Conditions</p> <p>The following conditions define how and when trusted entities can assume the role.</p> <p>There are no conditions associated with this role.</p>
--	--

Edit Trust Relationship

You can customize trust relationships by editing the following access control policy document.

Policy Document

```
1 {  
2   "Version": "2008-10-17",  
3   "Statement": [  
4     {  
5       "Sid": "",  
6       "Effect": "Allow",  
7       "Principal": {  
8         "Service": "ec2.amazonaws.com"  
9       },  
10      "Action": "sts:AssumeRole"  
11     },  
12     {  
13       "Effect": "Allow",  
14       "Principal": {  
15         "Service": "sagemaker.amazonaws.com"  
16       },  
17       "Action": "sts:AssumeRole"  
18     }  
19   ]  
20 }
```

Cancel

Update Trust Policy

Chapter 6: Analyzing Visitor Patterns to Make Recommendations

	poiID	poiName	lat	long	rideDuration	theme	theme2	theme3	theme4
0	1	Gadget's Go Coaster	33.810259	-117.918438	1.00	Kiddie	Roller Coaster	None	None
1	2	Astro Orbitor	28.418532	-81.579153	1.50	Spinning Ride	None	None	None
2	3	Mad Tea Party	33.813458	-117.918289	1.50	Family	Spinning Ride	None	None
3	4	Dumbo the Flying Elephant	33.813680	-117.918928	1.67	Family	Spinning Ride	None	None

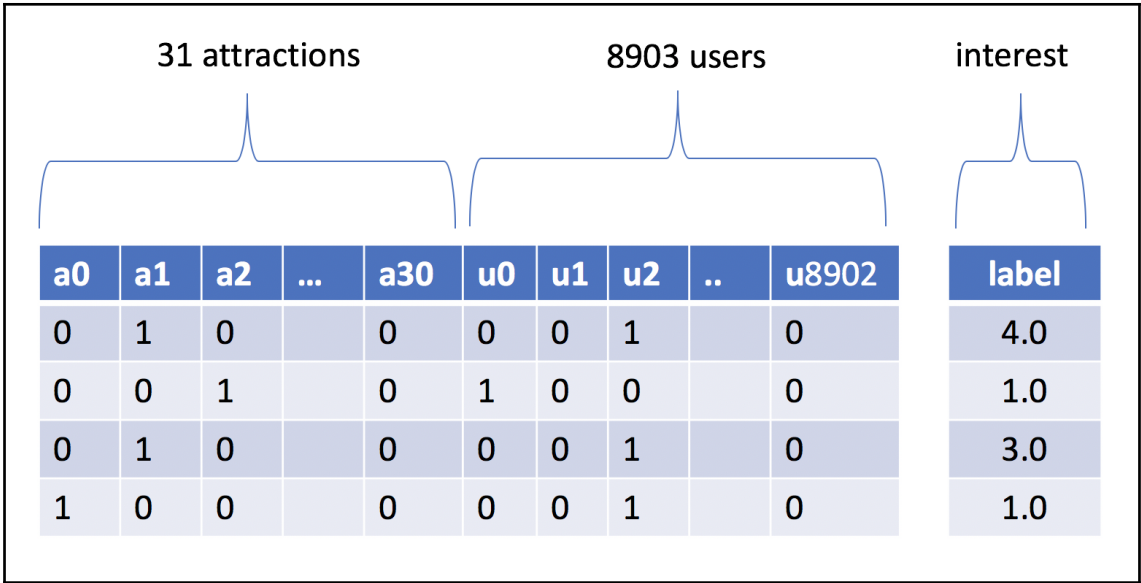
	id	nsid	takenUnix	poiID	poiTheme	poiFreq	rideDuration	seqID
0	5858403310	10004778@N07	1308262550	6	Ride	1665	120.0	1
1	5857850631	10004778@N07	1308270702	26	Family	18710	900.0	1
2	5858399220	10004778@N07	1308631356	6	Ride	1665	120.0	2
3	8277294024	10004778@N07	1355568624	26	Family	18710	900.0	3
4	9219062165	10004778@N07	1373030964	29	Water	10427	900.0	4

	summary	nsid	cnt
0	count	8903	8903
1	mean	None	4.86027181848815
2	stddev	None	5.965584836576787
3	min	10000151@N02	1
4	max	99987318@N03	31

	summary	nsid	poilD	count(1)
0	count	43271	43271	43271
1	mean	None	14.920061935245315	7.674678190936193
2	stddev	None	8.437883931275111	52.93100615991835
3	min	10000151@N02	1	1
4	max	99987318@N03	31	4128


	user_hash_id	poilD	pictures_taken
0	-1861435726	19	7
1	-1064654977	26	8
2	-636721096	17	1


	user_hash_id	recommendations
0	413285690	[(25, 39.260990142822266), (18, 34.83002853393...]
1	1005782960	[(29, 6.377601146697998), (25, 6.2345833778381...]
2	1410121870	[(25, 12.15351390838623), (29, 11.446855545043...]





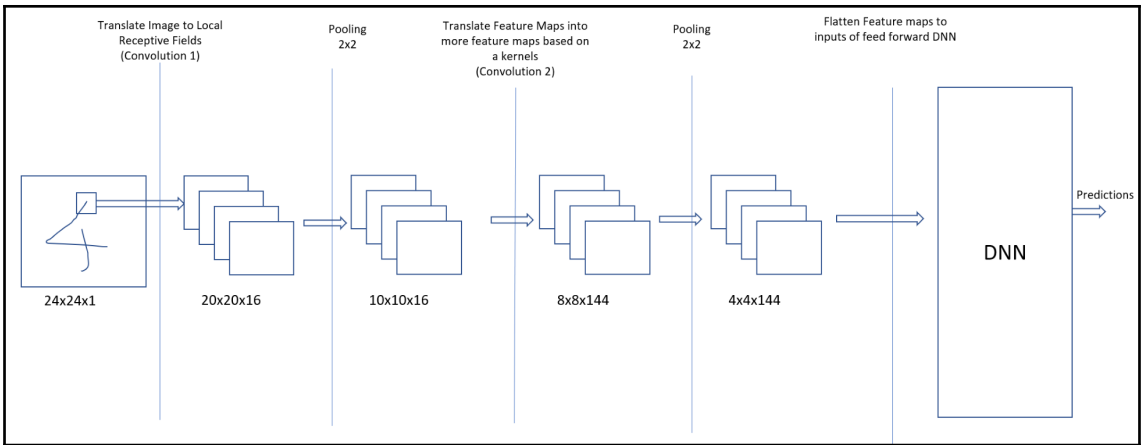

user_hash_id	poiID	pictures_taken	user_hash_id_index	user_hash_id_encoded	poi_id_indexed	poi_id_encoded	interest_level	features
-1861435726	19	7	279.0	(8903,[279],[1.0])	17.0	(31,[17],[1.0])	3.0	(8934,[17,310],[1....
-1064654977	26	8	181.0	(8903,[181],[1.0])	5.0	(31,[5],[1.0])	3.0	(8934,[5,212],[1....
-636721096	17	1	2187.0	(8903,[2187],[1.0])	4.0	(31,[4],[1.0])	1.0	(8934,[4,2218],[1....

Chapter 7: Implementing Deep Learning Algorithms

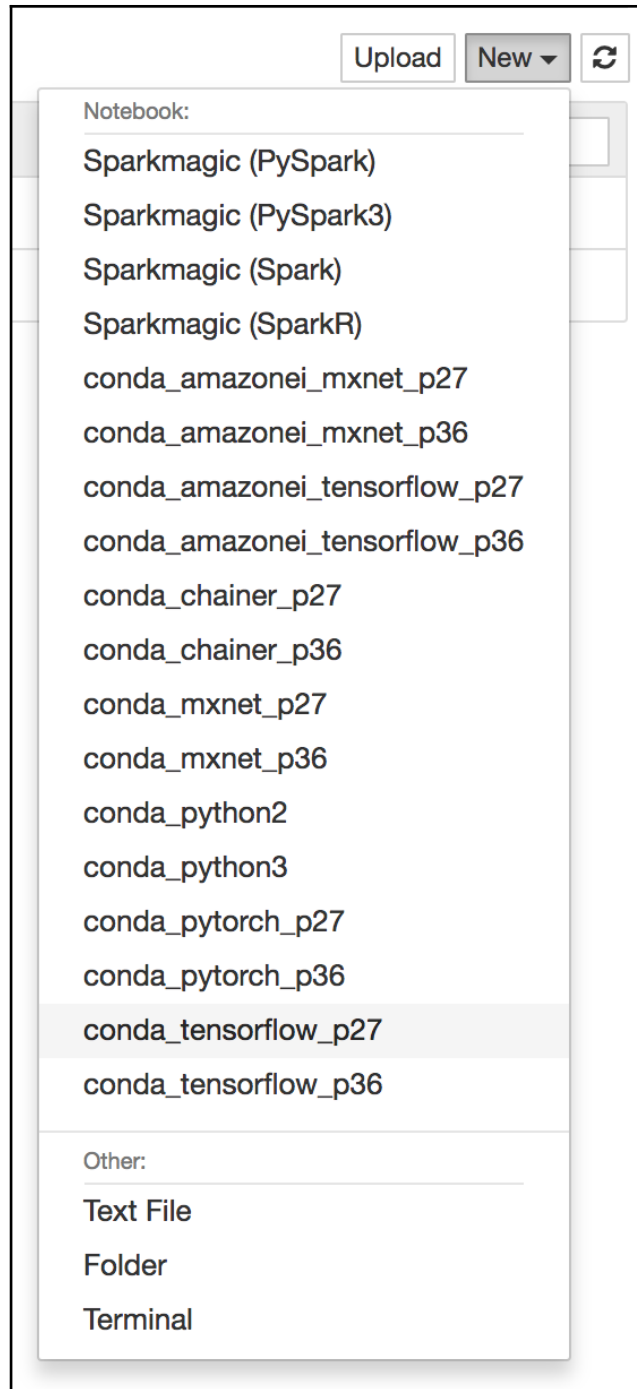









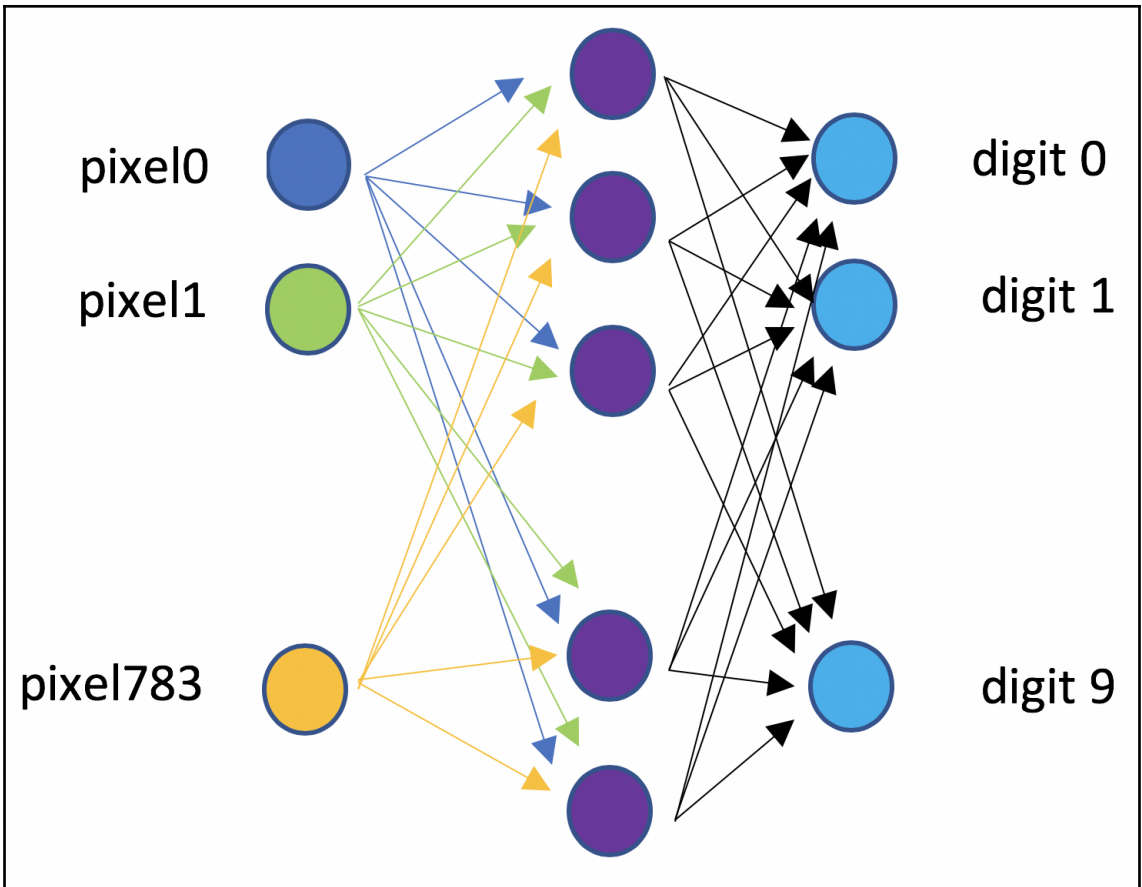
Chapter 8: Implementing Deep Learning with TensorFlow on AWS



	ID	crim	zn	indus	chas	nox	rm	age	dis	rad	tax	ptratio	black	lstat	medv	
	270	418	25.94060	0.0	18.10	0	0.679	5.304	89.1	1.6475	24	666	20.2	127.36	26.64	10.4
	205	309	0.49298	0.0	9.90	0	0.544	6.635	82.5	3.3175	4	304	18.4	396.90	4.54	22.8
	161	235	0.44791	0.0	6.20	1	0.507	6.726	66.5	3.6519	8	307	17.4	360.20	8.05	29.0



	label	pixel0	pixel1	pixel2	pixel3	pixel4	pixel5	pixel6	pixel7	pixel8	...	pixel774	pixel775	pixel776	pixel777	pixel778	pixel779	pixel780	pixel781	
0	1	0	0	0	0	0	0	0	0	0	...	0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0	0	0	...	0	0	0	0	0	0	0	0	0
2	1	0	0	0	0	0	0	0	0	0	...	0	0	0	0	0	0	0	0	0



```
show_image(0)
```

```
predicted digit: 6  
digit image:
```



```
show_image(3)
```

```
predicted digit: 2
```

```
digit image:
```



```
show_image(19)
```

```
predicted digit: 5  
digit image:
```

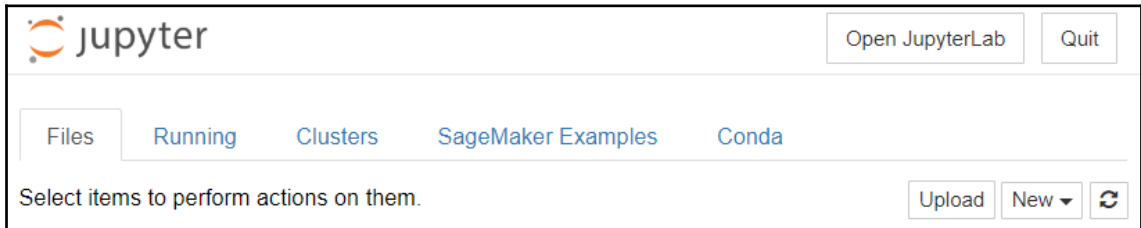


```
show_image(429)
```

```
predicted digit: 9  
digit image:
```



Chapter 9: Image Classification and Detection with SageMaker



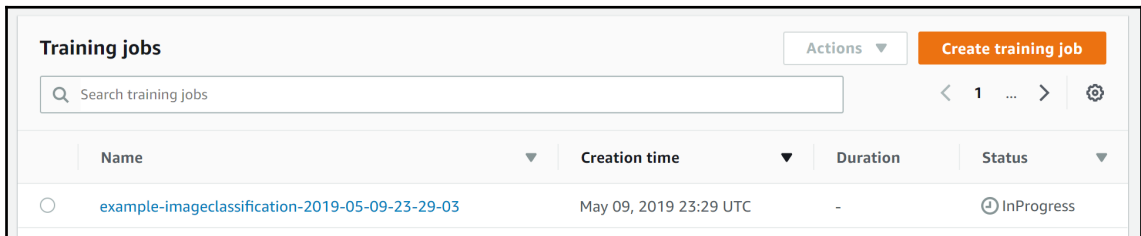
jupyter

Open JupyterLab Quit

Files Running Clusters SageMaker Examples Conda

Select items to perform actions on them.

Upload New ↕ ↻



Training jobs

Actions ▼ Create training job

Q Search training jobs < 1 ... > ⚙

	Name	Creation time	Duration	Status
○	example-imageclassification-2019-05-09-23-29-03	May 09, 2019 23:29 UTC	-	InProgress

Monitor

Access logs for debugging and progress reporting. View metrics to set alarms, send notifications, or take actions. [Learn more](#)

[View algorithm metrics](#)

[View logs](#)


[View instance metrics](#)

Add to dashboard


2019-05-09 (19:31:57) - 2019-05-09 (19:38:36)




DiskUtilization




MemoryUtilization




CPUUtilization




GPUUtilization




GPUMemoryUtilization




train:accuracy




validation:accuracy



train:accuracy:epoch



validation:accuracy:epoch



Chapter 10: Working with AWS Comprehend

Machine learning

Amazon Comprehend

Natural Language Processing and Text Analytics

Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to find insights and relationships in text.

Start Analyzing Text

We know you're curious. Jump in and try our APIs for Amazon Comprehend and Amazon Comprehend Medical.

[Launch Comprehend](#)

[Launch Comprehend Medical](#)

Entities | Key phrases | Language | Sentiment | Syntax

Analyzed text

Tim Cook traveled to New York for an Apple store opening.

▼ Results

Q Search < 1 > ⚙

Entity	Category	Confidence
Tim Cook	Person	0.99+
New York	Location	0.99+
Apple	Organization	0.96

Entities	Key phrases	Language	Sentiment	Syntax
Analyzed text				
I am not disappointed with switching to XXXX, this phone has a much better display and camera than the iPhones do. The pictures in the dark are phenomenal. It was easy to transfer all the data from my iPhone because of the easy set up features. If your thinking about switching from the iPhone, do it, you will not be disappointed.				
▼ Results				
Sentiment				
Neutral 0.00 confidence	Positive 0.98 confidence	Negative 0.00 confidence	Mixed 0.01 confidence	

Entities	Key phrases	Language	Sentiment	Syntax
Analyzed text				
Received the phone. It worked fine for the first 6 weeks,then I lost battery power I turned the phone off at night while charging, did not help. Then every else started to fail.				
▼ Results				
Sentiment				
Neutral 0.03 confidence	Positive 0.03 confidence	Negative 0.70 confidence	Mixed 0.22 confidence	

Classifier settings

Name

The name can be from 1 to 256 characters. Valid characters are a-z, A-Z, 0-9, period (.), colon (:), plus (+), equals (=), at sign (@), percent (%), and hyphen (-).

Language

Classifier encryption [Info](#)

Training data [Info](#)

Training document file must have one label and document per line.

Example

CAT	document text 1
CAT	document text 2
DOG	document text 3

File format: csv. Must be at least 1K documents per label.

S3 location

Search for a bucket or folder location in S3, or paste the URL of an input data file in S3

You should provide 1000 documents for each classification tag you're using. These should be in a single file in csv format, with one document per line. [Info](#)

Output data - optional [Info](#)

S3 location

Search for a bucket or folder location in S3, or paste the URL of a bucket or folder location in S3

Encryption [Info](#)

IAM role [Info](#)

Access permissions

Use an existing IAM role
 Create an IAM role

IAM role

Choose a role that grants access to the S3 input and output locations. Filtered by trust policy: comprehend.amazonaws.com

AmazonComprehendServiceRoleS3FullAccess-training ▼




▼ **Tags - optional**

A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter resources or track your AWS costs.

Key [Info](#) Value - optional [Info](#)

Classifiers

	Name	Start	Status
<input type="radio"/>	20NG	5/1/2019, 9:40:46 AM	🕒 Training

Classifier details			
Name 20NG	Document classifier arn arn:aws:comprehend:us-east-2:017198808824:document-classifier/20NG	Number of labels 21	Input data location s3://masteringmlsagemaker/comprehend/train.csv 
Status  Trained	Start 5/1/2019, 9:40:46 AM	Number of trained documents 10144	Language English
End 5/1/2019, 10:31:43 AM	Number of test documents 1127	Classifier encryption -	
Output			
Data location s3://masteringmlsagemaker/comprehend/017198808824-CLR-4033780ce75db10eac511242cd74d256/output/output.tar.gz 		Encryption -	
Classifier performance Info			
Accuracy 0.9095	Precision 0.8719	Recall 0.8676	F1 score 0.86

Create analysis job [Info](#)

Job settings

Name

The name can be up to 256 characters long. Valid characters are a-z, A-Z, 0-9, spaces and `_ . : / = + - % @`

Analysis type [Info](#)

Select classifier

A classifier organizes documents into categories or classes. Choose the recognizer you've created before or [Train new classifier](#) if you don't have one yet.

[View classifier details](#) 

Job encryption [Info](#)

Input data [Info](#)

S3 location

Search for a bucket or folder location in S3, or paste the URL of an input data file in S3

Input format - optional [Info](#)

Output data [Info](#)

S3 location


Search for a bucket or folder location in S3, or paste the URL of a bucket or folder location in S3

Encryption [Info](#)

Chapter 11: Using AWS Rekognition

Object and scene detection

Rekognition automatically labels objects, concepts and scenes in your images, and provides a confidence score.



Done with the demo?

[Learn more](#)

▼ Results


Boat	99.8 %
Vehicle	99.8 %
Transportation	99.8 %
Watercraft	99.2 %
Vessel	99.2 %
Ferry	95.6 %

[Show more](#)

► Request

Facial analysis

Get a complete analysis of facial attributes, including confidence scores.



Done with the demo?

[Learn more](#)

▼ Results



looks like a face	99.9 %
appears to be female	96.5 %
age range	26 - 43 years old
not smiling	96.8 %
not wearing glasses	99.9 %
not wearing sunglasses	99.9 %


[Show more](#)

► Request

► Response

Image moderation

Rekognition automatically detects explicit or suggestive adult content in your images, and provides confidence scores.



Done with the demo?

[Learn more](#)

▼ Results

Suggestive	98.3 %
Female Swimwear Or Underwear	98.3 %


► Request

▼ Response

```
{
  "ModerationLabels": [
    {
      "Confidence": 98.39189910888672,
      "Name": "Female Swimwear Or Underwear",
      "ParentName": "Suggestive"
    },
    {
      "Confidence": 98.39189910888672,
      "Name": "Suggestive",
      "ParentName": ""
    }
  ]
}
```

Celebrity recognition

Rekognition automatically recognizes celebrities in images and provides confidence scores.



Done with the demo?
[Learn more](#)

Results



Jeff Bezos
[Learn More](#)

Match confidence 100 %

Request

Response


Face comparison

Compare faces to see how closely they match based on a similarity percentage.

Reference face



Comparison faces



Done with the demo?
[Learn more](#)

Results



=



Similarity 92.7 %



≠



≠



Text in image

Rekognition automatically detects and extracts text in your images. [Learn More](#)



Done with the demo?
[Learn more](#)

▼ Results

US English only

| YOU | LOOKINP |
| ATT | MEE |

► Request

► Response

Chapter 12: Building Conversational Interfaces Using AWS Lex

Create your bot

Amazon Lex enables any developer to build conversational chatbots quickly and easily. With Amazon Lex, no deep learning expertise is necessary—you just specify the basic conversational flow directly from the console, and then Amazon Lex manages the dialogue and dynamically adjusts the response. To get started, you can choose one of the sample bots provided below or build a new custom bot from scratch.

CREATE YOUR OWN TRY A SAMPLE

Custom bot BookTrip OrderFlowers ScheduleAppointment

Bot name

Language

Output voice


Session timeout


IAM role [AWSServiceRoleForLexBots](#)
Automatically created on your behalf


COPPA Please indicate if your use of this bot is subject to the [Children's Online Privacy Protection Act \(COPPA\)](#). [Learn more](#)
 Yes No

▼ Sample utterances ⓘ








e.g. I would like to book a flight. 

I am hungry 

Can you order some food 

I need to order food 


▼ Slots ⓘ

Priority	Required	Name	Slot type	Version	Prompt	Settings
		<i>e.g. Location</i>	<i>e.g. AMAZON.US...</i>		<i>e.g. What city?</i>	
1.	<input checked="" type="checkbox"/>	RestaurantName	AMAZON.FoodEst...	Built-in	Where do you want to order food from?	 
2.	<input checked="" type="checkbox"/>	Foodname	AMAZON.Food	Built-in	What would you like to eat	 
3.	<input checked="" type="checkbox"/>	Time	AMAZON.TIME	Built-in	What time do you want your food to get	 


▼ Confirmation prompt ⓘ

Confirmation prompt

Confirm

{Time} is ok with the restaurant {RestaurantName} to deliver food. Do you want to to place the order? 


Cancel (if the user says "no")

Okay. I will cancel your order 


▼ Fulfillment ⓘ

AWS Lambda function Return parameters to client

▼ Response ⓘ

|| Message Custom Markup ⓘ 

One of these messages will be presented at random.



 [Add Message](#)

> Test bot (Latest)

✓ Ready. Build complete.

Can you order some food

Where do you want to order food from?

McDonalds

What would you like to eat

Spicy Chicken Sandwich

What time do you want your food to get here?

in 1 hour

12:08 is ok with the restaurant McDonalds to deliver food. Do you want to place the order?

yes


Intent OrderFood is ReadyForFulfillment:
Foodname:Spicy Chicken Sandwich
RestaurantName:McDonalds Time:12:08

Clear chat history




Chat with your bot...

Chapter 13: Creating Clusters on AWS



Create notebook instance

Amazon SageMaker provides pre-built fully managed notebook instances that run Jupyter notebooks. The notebook instances include example code for common model training and hosting exercises. [Learn more](#) 

Notebook instance settings

Notebook instance name

Maximum of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an AWS Region.

Notebook instance type

Elastic Inference [Learn more](#) 

▶ Additional configuration

Chapter 14: Optimizing Models in Spark and SageMaker

Best training job | **Training jobs** | Job configuration | Hyperparameter configuration | Tags

Training job status counter

Completed **10** | In Progress **0** | Stopped **0** | Failed **0** (Retryable: 0, Non-retryable: 0)

Training jobs View logs View instance metrics Stop Create model

Sorting by objective metric value will display only jobs that have metric values.

< 1 > ⚙️

	Name	Status	Objective metric value	Creation time	Training Duration
<input type="radio"/>	xgboost-190407-1532-010-ca7ffbfe	Completed	0.653780996799469	Apr 07, 2019 15:45 UTC	3 minute(s)
<input type="radio"/>	xgboost-190407-1532-009-83cdf8e2	Completed	0.6544309854507446	Apr 07, 2019 15:42 UTC	2 minute(s)
<input type="radio"/>	xgboost-190407-1532-008-8e3cddc3	Completed	0.6542659997940063	Apr 07, 2019 15:41 UTC	2 minute(s)
<input type="radio"/>	xgboost-190407-1532-007-68089fb8	Completed	0.6525049805641174	Apr 07, 2019 15:41 UTC	2 minute(s)
<input type="radio"/>	xgboost-190407-1532-006-5780c4f8	Completed	0.6500300168991089	Apr 07, 2019 15:37 UTC	2 minute(s)
<input type="radio"/>	xgboost-190407-1532-005-0e830ada	Completed	0.6545940041542053	Apr 07, 2019 15:37 UTC	3 minute(s)
<input type="radio"/>	xgboost-190407-1532-004-bc8ee28d	Completed	0.6323189735412598	Apr 07, 2019 15:37 UTC	3 minute(s)
<input type="radio"/>	xgboost-190407-1532-001-137764d2	Completed	0.6455860137939453	Apr 07, 2019 15:33 UTC	2 minute(s)
<input type="radio"/>	xgboost-190407-1532-002-45012b63	Completed	0.6471719741821289	Apr 07, 2019 15:33 UTC	2 minute(s)
<input type="radio"/>	xgboost-190407-1532-003-4a9d9782	Completed	0.6437379729881287	Apr 07, 2019 15:33 UTC	2 minute(s)

Chapter 15: Tuning Clusters for Machine Learning

Create Cluster - Advanced Options [Go to quick options](#)

Step 1: Software and Steps
Step 2: Hardware
Step 3: General Cluster Settings
Step 4: Security

Software Configuration

Release:

<input checked="" type="checkbox"/> Hadoop 2.8.5	<input type="checkbox"/> Zeppelin 0.8.1	<input type="checkbox"/> Livy 0.5.0
<input type="checkbox"/> JupyterHub 0.9.4	<input type="checkbox"/> Tez 0.9.1	<input type="checkbox"/> Flink 1.7.1
<input type="checkbox"/> Ganglia 3.7.2	<input type="checkbox"/> HBase 1.4.9	<input checked="" type="checkbox"/> Pig 0.17.0
<input checked="" type="checkbox"/> Hive 2.3.4	<input type="checkbox"/> Presto 0.215	<input type="checkbox"/> ZooKeeper 3.4.13
<input type="checkbox"/> MXNet 1.3.1	<input type="checkbox"/> Sqoop 1.4.7	<input type="checkbox"/> Mahout 0.13.0
<input checked="" type="checkbox"/> Hue 4.3.0	<input type="checkbox"/> Phoenix 4.14.1	<input type="checkbox"/> Oozie 5.1.0
<input type="checkbox"/> Spark 2.4.0	<input type="checkbox"/> HCatalog 2.3.4	<input type="checkbox"/> TensorFlow 1.12.0

AWS Glue Data Catalog settings (optional)
 Use for Hive table metadata

Edit software settings
 Enter configuration Load JSON from S3

`classification=config-file-name,properties=[myKey1=myValue1,myKey2=myValue2]`

Add steps (optional)
Step type:
 Auto-terminate cluster after the last step is completed

Add database

Database name

► Description and location (optional)

Database: default
 Filter tables and views...
Tables (1): house_prices
Views (0)

```

1 CREATE EXTERNAL TABLE IF NOT EXISTS house_prices(
2   crim float, zn float, indus float, chas float, nox float, rm float,
3   age float, dis float, tax float, ptratio float, lstat float, medv float )
4 ROW FORMAT DELIMITED
5   FIELDS TERMINATED BY ','
6 STORED AS TEXTFILE
7 LOCATION 's3://mastering-ml-aws/chapter3/linearmodels/train/'

```

Run query (Run time: 0.34 seconds, Data scanned: 0 KB)

Results: Query successful.

```

1 select * from house_prices limit 10

```

Run query (Run time: 1.75 seconds, Data scanned: 18.46 KB)

Results

	crim	zn	indus	chas	nox	rm	age	dis	tax	ptratio	lstat	medv	/
1	14.8	0.95577	0.0	8.14	0.0	0.538	6.047	88.8	4.4534	307.0	21.0	17.28	
2	23.3	0.0456	0.0	13.89	1.0	0.55	5.888	56.0	3.1121	276.0	16.4	13.51	
3	17.0	1.41385	0.0	19.58	1.0	0.871	6.129	96.0	1.7494	403.0	14.7	15.12	
4	24.3	0.33983	22.0	5.86	0.0	0.431	6.108	34.9	8.0555	330.0	19.1	9.16	
5	20.6	0.03306	0.0	5.19	0.0	0.515	6.059	37.3	4.8122	224.0	20.2	8.51	
6	34.9	0.0837	45.0	3.44	0.0	0.437	7.185	38.9	4.5667	398.0	15.2	5.39	
7	22.0	0.06899	0.0	25.65	0.0	0.581	5.87	69.7	2.2577	188.0	19.1	14.37	
8	20.0	6.80117	0.0	18.1	0.0	0.713	6.081	84.4	2.7175	666.0	20.2	14.7	
9	23.1	0.0187	85.0	4.15	0.0	0.429	6.516	27.7	8.5353	351.0	17.9	6.36	
10	22.8	0.49298	0.0	9.9	0.0	0.544	6.635	82.5	3.3175	304.0	18.4	4.54	

Cluster: SampleCluster Waiting Cluster ready after last step completed.

[Summary](#)
[Application history](#)
[Monitoring](#)
[Hardware](#)
[Configurations](#)
[Events](#)
[Steps](#)
[Bootstrap actions](#)

Connections: [Spark History Server](#), [JupyterHub](#), [Resource Manager](#) ... [\(View All\)](#)

Master public DNS: ec2-34-228-215-253.compute-1.amazonaws.com [SSH](#)

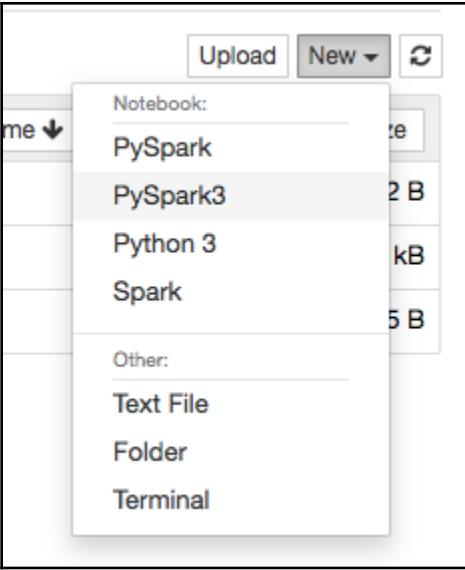
Tags: -- [View All / Edit](#)

Summary	Configuration details
ID: j-1CMVXX89792WV Creation date: 2019-04-11 13:42 (UTC-4) Elapsed time: 1 hour, 1 minute Auto-terminate: No Termination protection: On Change	Release label: emr-5.23.0 Hadoop distribution: Amazon 2.8.5 Applications: Spark 2.4.0, JupyterHub 0.9.4 Log URI: s3://aws-logs-095585830284-us-east-1/elasticmapreduce/ EMRFS consistent view: Disabled Custom AMI ID: --

Sign in

Username:

Password:



+
⌂
↶
↷
↕
↶
↷
▶ Run
■
↺
▶
Code
⌵
⌵

```
In [6]: %%sql
select * from house_prices limit 10
```

Type: Table Pie Scatter Line Area Bar

crim	zn	indus	chas	nox	rm	age	dis	tax	ptratio	lstat	medv
14.8	0.95577	0.0	8.14	0.0	0.538	6.047	88.8	4.4534	307.0	21.0	17.28
23.3	0.04560	0.0	13.89	1.0	0.550	5.888	56.0	3.1121	276.0	16.4	13.51
17.0	1.41385	0.0	19.58	1.0	0.871	6.129	96.0	1.7494	403.0	14.7	15.12
24.3	0.33983	22.0	5.86	0.0	0.431	6.108	34.9	8.0555	330.0	19.1	9.16
20.6	0.03306	0.0	5.19	0.0	0.515	6.059	37.3	4.8122	224.0	20.2	8.51
34.9	0.08370	45.0	3.44	0.0	0.437	7.185	38.9	4.5667	398.0	15.2	5.39
22.0	0.06899	0.0	25.65	0.0	0.581	5.870	69.7	2.2577	188.0	19.1	14.37
20.0	6.80117	0.0	18.10	0.0	0.713	6.081	84.4	2.7175	666.0	20.2	14.70
23.1	0.01870	85.0	4.15	0.0	0.429	6.516	27.7	8.5353	351.0	17.9	6.36
22.8	0.49298	0.0	9.90	0.0	0.544	6.635	82.5	3.3175	304.0	18.4	4.54

```
In [8]: df = spark.sql("select age,medv from house_prices")
```

```
In [10]: df.describe().show()
```

```

+-----+-----+
|summary|          age|          medv|
+-----+-----+
|  count|           266|           266|
|   mean|6.222233088392961|12.730187963722344|
| stddev|0.677306402813935| 7.124163055949489|
|    min|           3.561|           1.73|
|    max|           8.725|           37.97|
+-----+-----+

```

Chapter 16: Deploying Models Built in AWS

Endpoints					
<input type="text" value="Search endpoints"/>		Update endpoint		Actions	Create endpoint
<p style="text-align: right;">< 1 > ⚙️</p>					
Name	ARN	Creation time	Status	Last updated	
blazingtext-endpoint-2019-01-04-01	arn:aws:sagemaker:us-east-1:095585830284:endpoint/blazingtext-endpoint-2019-01-04-01	Apr 13, 2019 13:58 UTC	✔ InService	Apr 13, 2019 14:04 UTC	

blazingtext-endpoint-2019-01-04-01			Delete
Endpoint settings			
Name	Status	URL	
blazingtext-endpoint-2019-01-04-01	✔ InService	https://runtime.sagemaker.us-east-1.amazonaws.com/endpoints/blazingtext-endpoint-2019-01-04-01/invocations Learn more about the API	
ARN	Creation time	Last updated	
arn:aws:sagemaker:us-east-1:095585830284:endpoint/blazingtext-endpoint-2019-01-04-01	Sat Apr 13 2019 09:58:18 GMT-0400 (Eastern Daylight Time)	Sat Apr 13 2019 10:04:35 GMT-0400 (Eastern Daylight Time)	

Endpoint configuration settings							Change Clone
Name		Encryption key					
blazingtext-2019-01-04-01-42-36-379		-					
ARN		Creation time					
arn:aws:sagemaker:us-east-1:095585830284:endpoint-config/blazingtext-2019-01-04-01-42-36-379		Jan 04, 2019 01:48 UTC					
Production variants							
Model name	Training job	Variant name	Instance type	Elastic Inference	Initial instance count	Initial weight	
blazingtext-2019-01-04-01-48-23-529	blazingtext-2019-01-04-01-42-36-379	AllTraffic	mLm4.xlarge	-	1	1	

blazingtext-2019-01-04-01-48-23-529

Actions ▾

Create batch transform job

Create endpoint

Model settings

Name	ARN	Creation time	IAM role ARN
blazingtext-2019-01-04-01-48-23-529	arn:aws:sagemaker:us-east-1:095585830284:model/blazingtext-2019-01-04-01-48-23-529	Jan 04, 2019 01:48 UTC	arn:aws:iam::095585830284:role/service-role/AmazonSageMaker-ExecutionRole-20181129T145112 ↗

Container 1

Container Name	Model data location
Container 1	s3://mastering-ml-aws/chapter2/blazingtext/output/blazingtext-2019-01-04-01-42-36-379/output/model.tar.gz ↗
Image	
811284229777.dkr.ecr.us-east-1.amazonaws.com/blazingtext:latest	
Training job	
blazingtext-2019-01-04-01-42-36-379	

Create model

To deploy a model to Amazon SageMaker, first create the model by providing the location of the model artifacts and inference code. See [Deploying a Model on Amazon SageMaker Hosting Services](#) [Learn more about the API](#)

Model settings

Model name

Maximum of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an AWS Region.

IAM role

Amazon SageMaker requires permissions to call other services on your behalf. Choose a role or let us create a role that has the [AmazonSageMakerFullAccess](#) IAM policy attached.

Network

VPC - *optional*

For better security, we recommend that you use a private VPC.

Container definition 1

▼ Container input options

- Provide model artifacts and inference image location**
Use this for models trained using built-in algorithms, BYO algorithms, or models trained outside Amazon SageMaker.
- Use a model package resource**
Use this for model packages that contain inference images and artifacts from AWS Marketplace subscribed algorithms.
- Use a model package subscription from AWS Marketplace**
Use this for model packages published by vendors from AWS Marketplace.

▼ Provide model artifacts and inference image

Location of inference code image

Type the registry path where the inference code image is stored in Amazon ECR.

Location of model artifacts - *optional*

Type the URL where model artifacts are stored in S3.

The path must point to a single gzip compressed tar archive (.tar.gz suffix).

Container host name - *optional*

Type the DNS host name for the container.

Maximum of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an AWS Region.