

MLlib Tutorial

IPython notebooks consist of multiple "cells", which can contain text (markdown) or code. You can run a single cell, multiple cells at a time, or every cell in the file. IPython also gives you the ability to export the notebook into different formats, such as a python file, a PDF file, an HTML file, etc.

Spend a few minutes and make yourself familiar with the IPython interface, it should be pretty straight forward.

The Machine Learning Library contains common machine learning algorithms and utilities. A summary of these features can be found [here \(http://spark.apache.org/docs/latest/mllib-guide.html\)](http://spark.apache.org/docs/latest/mllib-guide.html).

In this tutorial, we will explore [collaborative filtering \(http://spark.apache.org/docs/latest/mllib-collaborative-filtering.html\)](http://spark.apache.org/docs/latest/mllib-collaborative-filtering.html) on a small example problem. For your submission, you must turn in a PDF of this IPython notebook, but fully completed with any missing code.

Collaborative Filtering Example

We will go through an exercise based on the example from the collaborative filtering page. The first step is to import the MLlib library and whatever functions we want with this module.

```
In [1]: from pyspark.mllib.recommendation import ALS, Rating
```

Loading the Input Data

Next, we must load the data. There are some example datasets that come with Spark by default. Example data related to machine learning in particular is located in the `$SPARK_HOME/data/mllib` directory. For this part, we will be working with the `$SPARK_HOME/data/mllib/als/test.data` file. This is a small dataset, so it is easy to see what is happening.

```
In [2]: data = sc.textFile("../data/mllib/als/test.data")
```

Even though, we have the environment `$SPARK_HOME` defined, but it can't be used here. You must specify the full path, or the relative path based off where the IPython file is located.

The `textFile` command will create an RDD where each element is a line of the input file. In the below cell, write some code to (1) print the number of elements and (2) print the fifth element. Print your result in a single line with the format: "There are X elements. The fifth element is: Y".

```
In [3]: print "There are %d elements. The fifth element is: %s."%(data.count(),data.collect()[4])
```

There are 16 elements. The fifth element is: 2,1,5.0.

Transforming the Input Data

This data isn't in a great format, since each element in the RDD is currently a string. However, we will assume that the first column of the string represents a user ID, the second column represents a product ID, and the third column represents a user-specified rating of that product.

In the below cell, write a function that takes a string (that has the same format as lines in this file) as input and returns a tuple where the first and second elements are ints and the third element is a float.

Call your function parser.

We will then use this function to transform the RDD.

```
In [4]: import re
def parser(inn):
    out=[]
    out.append(re.search('(.),(.+),(.+)',inn).group(1))
    out.append(re.search('(.),(.+),(.+)',inn).group(2))
    out.append(re.search('(.),(.+),(.+)',inn).group(3))
    return out
```

```
In [5]: ratings = data.map(parser).map(lambda l: Rating(*l))
ratings.collect()
```

```
Out[5]: [Rating(user=1, product=1, rating=5.0),
Rating(user=1, product=2, rating=1.0),
Rating(user=1, product=3, rating=5.0),
Rating(user=1, product=4, rating=1.0),
Rating(user=2, product=1, rating=5.0),
Rating(user=2, product=2, rating=1.0),
Rating(user=2, product=3, rating=5.0),
Rating(user=2, product=4, rating=1.0),
Rating(user=3, product=1, rating=1.0),
Rating(user=3, product=2, rating=5.0),
Rating(user=3, product=3, rating=1.0),
Rating(user=3, product=4, rating=5.0),
Rating(user=4, product=1, rating=1.0),
Rating(user=4, product=2, rating=5.0),
Rating(user=4, product=3, rating=1.0),
Rating(user=4, product=4, rating=5.0)]
```

Your output should look like the following:

```
[Rating(user=1, product=1, rating=5.0),
Rating(user=1, product=2, rating=1.0),
Rating(user=1, product=3, rating=5.0),
Rating(user=1, product=4, rating=1.0),
Rating(user=2, product=1, rating=5.0),
Rating(user=2, product=2, rating=1.0),
Rating(user=2, product=3, rating=5.0),
Rating(user=2, product=4, rating=1.0),
Rating(user=3, product=1, rating=1.0),
Rating(user=3, product=2, rating=5.0),
Rating(user=3, product=3, rating=1.0),
Rating(user=3, product=4, rating=5.0),
Rating(user=4, product=1, rating=1.0),
Rating(user=4, product=2, rating=5.0),
Rating(user=4, product=3, rating=1.0),
Rating(user=4, product=4, rating=5.0)]
```

If it doesn't, then you did something wrong! If it does match, then you are ready to move to the next step.

Building and Running the Model

Now we are ready to build the actual recommendation model using the Alternating Least Squares algorithm. The documentation can be found [here](http://spark.apache.org/docs/1.5.2/api/python/pyspark.mllib.html#module-pyspark.mllib.recommendation) (<http://spark.apache.org/docs/1.5.2/api/python/pyspark.mllib.html#module-pyspark.mllib.recommendation>), and the papers the algorithm is based on are linked off the [collaborative filtering page](http://spark.apache.org/docs/latest/mllib-collaborative-filtering.html) (<http://spark.apache.org/docs/latest/mllib-collaborative-filtering.html>).

```
In [6]: rank = 10
        numIterations = 10
        model = ALS.train(ratings, rank, numIterations)
```

```
In [7]: # Let's define some test data
        testdata = ratings.map(lambda p: (p[0], p[1]))

        # Running the model on all possible user->product predictions
        predictions = model.predictAll(testdata)
        local=predictions.collect()
```

Transforming the Model Output

This result is not really in a nice format. Write some code that will transform the RDD so that each element is a user ID and a dictionary of product->rating pairs. Note that for the a Ratings object (which is what the elements of the RDD are), you can access the different fields by via the `.user`, `.product`, and `.rating` variables. For example, `predictions.take(1)[0].user`.

Call the new RDD `userPredictions`. It should look as follows (when using `userPredictions.collect()`):

```
[(4,
  {1: 1.0011434289237737,
   2: 4.996713610813412,
   3: 1.0011434289237737,
   4: 4.996713610813412}),
 (1,
  {1: 4.996411869659315,
   2: 1.0012037253934976,
   3: 4.996411869659315,
   4: 1.0012037253934976}),
 (2,
  {1: 4.996411869659315,
   2: 1.0012037253934976,
   3: 4.996411869659315,
   4: 1.0012037253934976}),
 (3,
  {1: 1.0011434289237737,
   2: 4.996713610813412,
   3: 1.0011434289237737,
   4: 4.996713610813412})]
```

```
In [8]: users=[4,1,2,3]
dicts=[]
for i in range(4):
    dicts.append({})
for i in local:
    for j in range(4):
        if users[j]==i.user:
            dicts[j][i.product]=i.rating
userPredictions=sc.parallelize(zip(users,dicts))
userPredictions.collect()
```

```
Out[8]: [(4,
  {1: 1.0008106721056471,
   2: 4.9957732449084755,
   3: 1.0008106721056471,
   4: 4.9957732449084755}),
 (1,
  {1: 4.9966717169964285,
   2: 1.0006309210451323,
   3: 4.9966717169964285,
   4: 1.0006309210451323}),
 (2,
  {1: 4.9966717169964285,
   2: 1.0006309210451323,
   3: 4.9966717169964285,
   4: 1.0006309210451323}),
 (3,
  {1: 1.0008106721056471,
   2: 4.9957732449084755,
   3: 1.0008106721056471,
   4: 4.9957732449084755})]
```

Evaluating the Model

Now, lets calculate the mean squared error.

```
In [9]: userPredictions = predictions.map(lambda r: ((r[0],r[1]), r[2]))
ratesAndPreds = ratings.map(lambda r: ((r[0],r[1]), r[2])).join(userPredictions)
MSE = ratesAndPreds.map(lambda r: (r[1][0] - r[1][1])**2).mean()
print("Mean Squared Error = " + str(MSE))
```

Mean Squared Error = 0.0

Reflections

Why do you believe that this model achieved such a low error. Is there anything we did incorrectly for testing this model?

We specified the value of every customer and every product. Therefore we can have perfect prediction...since we already have all the actual values of anything we might try to predict. I think the proper turn of phrase here might be "hindsight is 20/20"