

HOW TO RUN MACHINE LEARNING TRANSFORMS IN AWS GLUE



Here we show you how to do a machine learning transformation with Amazon Glue. Previous Glue tutorials include:

- [How To Make a Crawler in Amazon Glue](#)
- [How To Join Tables in Amazon Glue](#)
- [How To Define and Run a Job in AWS Glue](#)
- [AWS Glue ETL Transformations](#)

Now, let's get started.

Amazon's machine learning

A fully managed service from Amazon, AWS Glue handles data operations like ETL to get your data prepared and loaded for analytics activities. Glue can crawl S3, DynamoDB, and JDBC data sources.

Amazon called their offering **machine learning**, but they only have one ML-type function, **findMatches**. It uses an ML algorithm, but Amazon does not tell you which one. They even boast on their web page you don't need to know—but a [data scientist](#) would certainly want to know.

You can study their execution log to gain some insight into what their code is doing. Suffice it to say it is doing a type of [clustering algorithm and using Apache Spark](#) as a platform to execute that.

The process: Amazon Glue machine learning

Here is the general process for running machine learning transformations:

1. Upload a csv file to an S3 bucket. Then you [set up a crawler](#) to crawl all the files in the designated S3 bucket. For each file it finds, it will create a metadata (i.e., schema) file in Glue that contains the column names.
2. Set up a **FindMatches** machine learning task in Glue. It's an iterative process. It takes your input data, created in the crawler process, and makes a label file. These labels are like a k-means clustering algorithm. It looks at the input data and all of the columns in the data set. Then it put the data into groups, each labeled with a **labeling_set_id**.
3. Download the label file. There will be an empty column called **label**. You are invited to add your own label to classify data however you see fit. For example, it could be borrower risk rating, whether or not a patient has diabetes, or whatever. Labels should be a single value, like A, B, C or 1, 2, 3. A data scientist would say they must be **categorical**.
4. Upload the labelled file to a different S3 bucket. Do not use the same bucket where you put the original input data, as the crawler will attempt to crawl that and create another metadata file.
5. Rerun Step 2, above, and it creates another labelled file. Do this iteratively until it supplies the most accurate result. In this example, there was no improvement from one run to the next. Repeating machine learning runs is standard practice for improving accuracy. However, at some point, the gain in accuracy will level off.
6. Generate and then inspect the **Quality Metrics**. Perhaps change some of the parameters and run the **Tune** operation, which means to run the algorithm again.

Tutorial: Amazon Glue machine learning

Now, let's run an example to show you how it works.

I have copied the Pima Native American database [from Kaggle](#) and put it on GitHub, [here](#). You have to add a primary key column to that data, which Glue requires. Download the data [here](#). I have also copied the input data and the first and second label files [here](#), in a Google Sheet, so that you can see the before and after process.

The data looks like this:

```
recordID,Pregnancies,Glucose,BloodPressure,SkinThickness,Insulin,BMI,Diabetes
PedigreeFunction,Age,Outcome
2,6,148,72,35,0,33.6,0.627,50,1
3,1,85,66,29,0,26.6,0.351,31,0
4,8,183,64,0,0,23.3,0.672,32,1
5,1,89,66,23,94,28.1,0.167,21,0
6,0,137,40,35,168,43.1,2.288,33,1
```

Then copy it to an Amazon S3 bucket as shown below. You need to have installed the Amazon CLI (command line interface) and run **aws configure** to configure your credentials. Importantly, the data must be in the same [Amazon zone](#) as the instance you are logged into.

```
aws s3 cp diabetes.csv s3://sagemakerwalkerm1
```

Add a label

The diabetes data is already labelled in the column **outcome**. So, I used Google Sheets to copy that value into the label **column**. You do this after the first run, like this:

1. Upload the original data.
2. Run a training model
3. Download the resulting labels file.

At that point you can populate the label with some kind of categorical data. You might put the outcome of logistic regression on your input data set into this label, but that's optional. You don't need a label at all.

The algorithm does not require a label the first time it runs. Glue says:

As you can see, the scope of the labels is limited to the `labeling_set_id`. So, labels do not cross `labeling_set_id` boundaries.

In other words, when there is no label, it groups records by **labeling_set_id** without regards to the **label** value. When there is a label then the **labeling_set_id** is within the label.

In other words, given this:

labeling_set_id	labeling_set_id	other columns
123	blank	
123	blank	
456	blank	

The first two rows are grouped together. But if we add a label:

labeling_set_id	label	other columns
123	A	
123	A	
456	B	
456	B	
789	A	

Then the first two rows are matched together while rows 456, even if they had matched without the label, are groups separately. Remember Amazon said they "don't cross boundaries."

Of course, that does not mean only the label determines what is considered to be a match. (That would be of little use.) It's the other columns that determine what matches. The label just confines that matching to records with that label. So, it's matching within a subset of records. It's like having n number of files with no label instead of one file with n labels, so you can run the process one time and not n times.

Anyway, that's the conclusion I draw from this design. Perhaps yours will differ.

Crawl S3

We start with the crawlers. Here is the metadata extracted from the diabetes.csv file in S3:

Schema			
	Column name	Data type	Partition key
1	recordid	bigint	
2	pregnancies	bigint	
3	glucose	bigint	
4	bloodpressure	bigint	
5	skinthickness	bigint	
6	insulin	bigint	
7	bmi	double	
8	diabetespedigreefunction	double	
9	age	bigint	
10	outcome	bigint	

It created these tables in the database.

Choose a data source				
AWS Glue data catalog				
<input type="text"/> Filter by attributes or search by keyword				
Name	Database	Location	Classi	
<input checked="" type="radio"/> diabetes_csv	diabetes	s3://sagemakerwalkerm1/diabetes.csv	csv	
<input type="radio"/> s3_	diabetes	s3://sagemakerwalkerm1/s3/	Unkno	
<input type="radio"/> sagemaker	diabetes	s3://sagemakerwalkerm1/sagemaker/	Unkno	

Pick an IAM role that has access to S3 and give the transformation a name.

Configure transform properties

Transform name

Description (optional)

IAM role ⓘ

AWSGlueServiceRole-S3IAMRole
▼
↻

Ensure that this role has permission to your Amazon S3 sources, temporary directory, and labeling files. [Create IAM role](#)

▶ Task run properties (optional)

▶ Tags (optional)

Next

The data must have a primary key. The matching algorithm requires that to do its matching logic.

Choose a primary key

Choose a primary key column. This column typically contains a unique identifier for every record in the data source.

	Column name	Data type
<input checked="" type="radio"/>	recordid	bigint
<input type="radio"/>	pregnancies	bigint
<input type="radio"/>	glucose	bigint

Showing: 1 -

Then it asks

you to tune the transformation. These are tradeoffs between cost and accuracy:

- **Cost** is financial.
- **Cost function** is data science and computing.

(Pricing is based on resources (DPUs) you consume, which [I cover below.](#))

The data science-related tuning parameters are between **recall** and **precision**.

Recall is:

$$\frac{\text{true positives}}{\text{predicted results}} = \frac{\text{true positives}}{\text{true positives} + \text{false positives}}$$

Precision is:

$$\frac{\text{true positives}}{\text{actual results}} = \frac{\text{true positives}}{\text{true positive} + \text{false negative}}$$

Tune transform

Select a tuning option or adjust the sliders to tune the transform. Choosing a midpoint optimizes the transform to find duplicates in a reasonable time.

Recall vs. precision



- Balanced (0.5)** Even tradeoff between recall and precision
- Favor recall (0.2)** Find more matches even if some are incorrectly matched
- Favor precision (0.9)** Find matches with fewer incorrect matches
- Custom

The value must be between 0.0 and 1.0 (inclusive).

Lower cost vs. accuracy



- Balanced (0.5)** Even tradeoff between accuracy and cost
- Favor lower cost (0.2)** Use less resources to find matches
- Favor accuracy (0.9)** Use more resources but potentially find more matches
- Custom

Here is a summary of the parameters:

Review

Machine learning transforms

Transform name	diabetes
IAM role	arn:aws:iam::782976337272:role/service-role/AWSGlueServiceRole-S3IAMRole
Data source	diabetes_csv
Primary key	recordid
Precision-recall tradeoff	0.5
Accuracy-cost tradeoff	0.5
Force output to match labels	false
Tags	-

Task run properties

Worker type	G.2X
Number of workers	10
Task timeout (minutes)	2880
Number of retries	0
Spark Version	2.4

[Back](#)

[Finish](#)

- **The first time** let it generate a label file for you. It will match records based on all of the data points taken together.
- **The second time** it will incorporate labels in its matching algorithm should you choose to add one.

Teach the transform using labels

Teach your machine learning algorithms by providing examples, called labels. For your transform, provide examples of matching and non-matching records.

Labeling file

- I do not have labels
 I have labels

Generate the labeling file

AWS Glue extracts records from your source data and suggests potential matching records. The file will contain approximately 100 data samples for you to work with.

Generate labeling file

You can download the file once it has been generated.

Download labeling file

Next

Edit the file in Excel or Google Sheets to both review it and optionally add a label. Copy it back to S3, putting it in a different bucket than the original upload file. Then run transformation again (called **train**). It will produce yet another label file which is the results of the matching aka grouping process.

Upload labels

Labeling file

- I do not have labels
 I have labels

Upload labels from S3

The completed labeling file must be in the correct format and in Amazon S3.

Upload labeling file from S3

Back

Next

It asks for the bucket name:

Generate labeling file

S3 path to store the generated label file

s3://bucket/prefix/object

Generate

You download the labels from this screen.

tfm-e52040b35753cb0478fd5895acd0691de34a5ada_la...

Overview

Properties

Permissions

Select from

Open

Download

Download as

Make public

Copy path

Owner

critique_american

Last modified

Aug 27, 2020 2:16:55 PM GMT+0300

Etag

3f39123fa844449f3182380382b64c30

Storage class

Standard

Server-side encryption

Here is the first label

file it created. You can't see all of the columns because it's too wide. But you can see the labeling_set_id, thus how it grouped the data:

	A	B	C	D	E	F	G	H	I	
1	labeling_set_id	label	pregnancies	bmi	outcome	age	diabetespedigre	bloodpressure	skinthickness	rec
2	244fcf87-d3a6-36cb-81be-43a29dece209	0	0	40.5	0	44	1.781	82	0	
3	244fcf87-d3a6-36cb-81be-43a29dece209	0	1	22.4	0	27	0.207	122	0	
4	244fcf87-d3a6-36cb-81be-43a29dece209	1	10	32.4	1	42	0.272	66	0	
5	244fcf87-d3a6-36cb-81be-43a29dece209	0	1	24.3	0	21	0.187	58	0	
6	244fcf87-d3a6-36cb-81be-43a29dece209	0	0	18.4	0	27	0.582	76	0	
7	244fcf87-d3a6-36cb-81be-43a29dece209	0	2	25.3	0	22	0.881	62	10	
8	244fcf87-d3a6-36cb-81be-43a29dece209	0	6	30.8	0	37	0.122	70	32	
9	244fcf87-d3a6-36cb-81be-43a29dece209	0	1	32	0	22	0.389	68	35	
10	244fcf87-d3a6-36cb-81be-43a29dece209	0	0	20	0	22	0.236	68	22	
11	244fcf87-d3a6-36cb-81be-43a29dece209	1	13	42.3	1	44	0.257	114	0	
12	28603e1d-6ae9-38af-9801-c5bb2d7ff7c3	0	7	27.4	0	40	0.294	64	0	
13	28603e1d-6ae9-38af-9801-c5bb2d7ff7c3	1	7	30.4	1	36	0.383	66	0	
14	28603e1d-6ae9-38af-9801-c5bb2d7ff7c3	0	2	28.7	0	25	0.092	68	22	
15	28603e1d-6ae9-38af-9801-c5bb2d7ff7c3	1	1	43.3	1	41	0.282	0	0	
16	28603e1d-6ae9-38af-9801-c5bb2d7ff7c3	0	0	22.1	0	21	0.207	62	17	
17	28603e1d-6ae9-38af-9801-c5bb2d7ff7c3	1	3	31.6	1	28	0.268	70	18	
18	28603e1d-6ae9-38af-9801-c5bb2d7ff7c3	0	3	26.3	0	24	0.107	74	15	
19	28603e1d-6ae9-38af-9801-c5bb2d7ff7c3	0	0	36.3	0	23	0.804	80	37	

Evaluation metrics

This screen lets you calculate accuracy. I have yet to figure out where you can see the results as the screen mentioned in the documentation does not exist. (I will update this tutorial once I get a response on the user forum.)

Estimate quality metrics (optional)

Estimate your transform's ability to find matches. Estimates are calculated by comparing the transform match predictions using a subset of your labeled data against the labels you have provided. These estimates are approximate. To improve your transform quality, provide more labels.

New estimates start a task that you can monitor in the **History** pane of the transform. When the task completes, new estimates can be viewed in the **Quality metrics** pane of the transform.

[Estimate transform quality](#)

i You have not yet estimated the quality of the transform. Click "Estimate transform quality" to get the quality of your transform.

Quality metric	Definition	Result	Last modified
Area under the Precision-Recall curve	Single number summarizing the performance of the transform	-	-
Precision	When your transform predicts a match, how often is it correct?	-	-
Recall upper limit	For an actual match, how often does your transform predict a match?	-	-
F1	Indicates transform's accuracy. Harmonic mean of Precision and Recall.	-	-

* Metrics shown are from the last quality estimation run.

** End-to-End recall will tend to be closer to the upper limit as the cost-accuracy slider favors accuracy. See documentation for additional information about End-to-End recall.

Pricing

Price is by DPU. I used 10 DPUs for about 30 minutes. It's \$0.44 for each multiple or fraction of an hour. So presumably I spent $\$0.44 \times 10 = \4.40 .

These ETL jobs run on Amazon's Spark and Yarn infrastructure. If you want to write code to do transformations you need to set up a **Development Endpoint**. Basically, the development endpoint is a [virtual machine](#) configured to run Spark and Glue. We explained how to use a Development Endpoint [here](#). Then you can run Python or Scala and optionally [use a Jupyter Notebook](#).

Important note: When you don't need your development endpoint, be sure to delete it—it gets expensive quickly! (I spent \$1,200 on that in a month.)

Additional resources

For more tutorials like this, explore these resources:

- [BMC Machine Learning & Big Data Blog](#)
- [Apache Spark Guide](#), with 15 articles and tutorials
- [AWS Guide](#)
- [Amazon Braket Quantum Computing: How To Get Started](#)