

INTRODUCTION TO GOOGLE CLOUD MACHINE LEARNING ENGINE



The Google Cloud Machine Learning Engine is almost exactly the same as [Amazon Sagemaker](#). It is not a SaaS program that you can just upload data to and start using like the [Google Natural Language API](#). Instead, you have to program Google Cloud ML using any of the ML frameworks such as [TensorFlow](#), [scikit-learn](#), [XGBoost](#), or [Keras](#). Then Google spins up an environment to run the training models across its cloud. So in a word, GC ML is just a computing platform to run ML jobs. By itself it does not do ML. You have to code that.

For example, you use the command below to start or otherwise manage a job. This is the command line interface to the Google Cloud ML Engine [SDK](#). You install that and other SDKs locally on a VM to work with the product.

gcloud ml-engine GROUP | COMMAND

What these jobs do is train algorithms and then make predictions or do classification. That is what ML does. Doing this involves solving very large sets of equations at the same time, which requires lots of computing power. So what the GC ML platform does for you is provision the resources needed to do that. Amazon SageMaker uses containers. Google says they do not. So they use some other container-like approach. (They could also be providing you with access to Google [TPUs](#), which are Google proprietary application-specific chips tuned for ML mathematics.)

In other words, most ML works on the same principle which is to iteratively look at a set of equations

and then try different coefficients to minimize the error using some loss function. This requires multiplying n-dimensional matrices which is a problem that lends itself well to running across a cloud. This is because such a problem can be divided into smaller pieces.

GC ML creates this cloud on-the-fly for you. So you do not have to set up virtual machines and containers ahead of time by yourself. That is the value that it delivers.

It can take some time, sometimes hours or even days, to train a neural network or other ML problem. But once saved that data model be used to do classification or make predictions. So you do not have to run that large training job again.

Google even says it will let you import training models created on other platforms. I did not look at that part but presumably that must let you import, for example, training models saved as Hadoop Parquet files on another cloud.

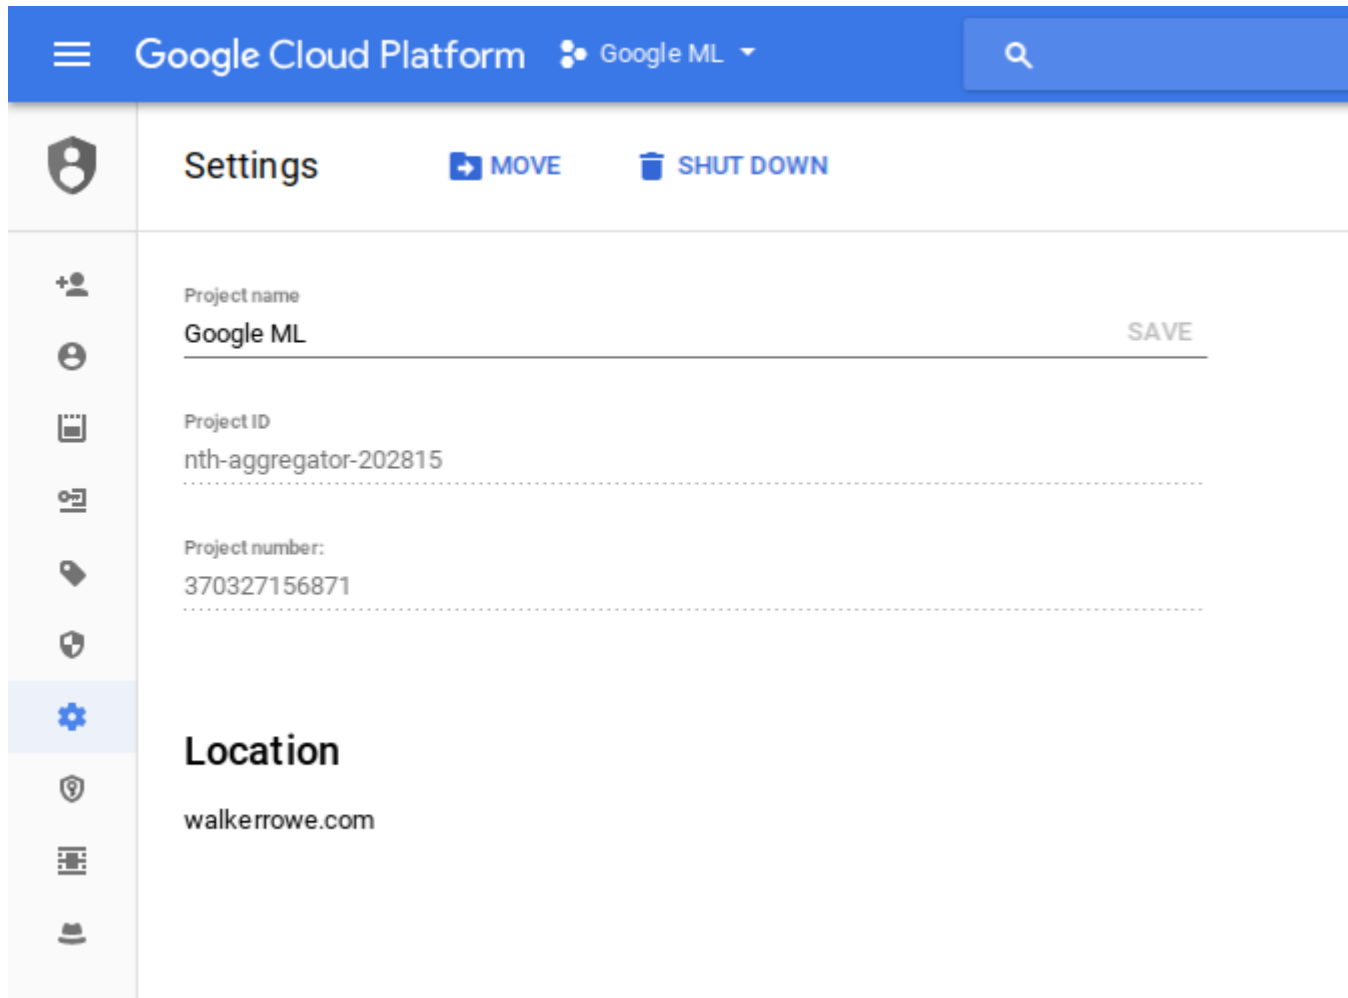
Getting Started

To get started with the product you can walk through [this tutorial](#) by Google.

This example is Python code using TensorFlow. To use it you:

- provision a Google Cloud virtual machine (so you can have some place to write the code).
- install the Google Cloud SDKs.
- create a model in the Google Cloud ML console.
- run the `gcloud ml-engine` command (shown below) to kick off the job.
- inspect the output.

For example, the console looks like this:



You

submit a job like this:

```
gcloud ml-engine jobs submit training $JOB_NAME \
--job-dir $OUTPUT_PATH \
--runtime-version 1.4 \
--module-name trainer.task \
--package-path trainer/ \
--region $REGION \
-- \
--train-files $TRAIN_DATA \
--eval-files $EVAL_DATA \
--train-steps 1000 \
--eval-steps 100 \
--verbosity DEBUG
```

This command-line approach is different from Amazon SageMaker which uses a graphical interface.

Pricing

Google Cloud ML is not free. However, when you first sign up you get \$300 credit that is good for 1 year. Contrast that with Amazon, where you start incurring charges right away. Charges from both companies coming from the provisioning of virtual machines and containers needed to run the ML jobs. So they do not charge you as you write code. They just charge you when you run it.

Required Skills

All of the ML frameworks on the Google ML cloud are written in Python. These are all open source projects.

They are all complicated and require an understanding of data science to use. That means you need to know linear algebra, advanced statistics, neural networks, gradient descent, regression, etc. So these are not tools for ordinary programmers, they are for data scientists.

What Can you do with Google Cloud ML

Google Cloud ML lets you solve business and scientific problems such as logistic and linear regression, classification, and neural networks. You can use these to make predictions and to classify data. For example, it could do handwriting recognition. But a more business-like problem would be something like vehicle preventive maintenance or calculating the efficacy of the advertising budget.