

Churn rate prediction framework

You can predict how likely a customer is to churn based on how recently they've used your service. But unless you have a very small number of customers, spending a huge amount of time manually drilling down into the data and activity of each customer just isn't viable.

However, there is a solution - and it lies in utilizing machine learning.

Machine learning is an AI-based data analysis technique.

Machine learning programs incorporate 'models'.

An AI model is a program or algorithm that uses a set of data to recognize certain patterns.

To enable the model to do this you need to feed it as much data as possible and, once it has this data, it can apply what it's learned to any new data set you give it, and can predict possible future actions based on past patterns.

So let's look at your first step:

Step 1 - Collecting data

First off, you need the data to feed to your model. This includes existing customers and customers who have already churned. Each piece of customer information you collect is called a feature.

The more features you can collect, the more accurate your model will be.

There are four main areas you need to focus on:

Customer features - This is data directly related to the characteristics of your customers, so everything from age and gender to education and income level.

Support features - This is information related to how customers interact with your customer support. How often they contact customer support and the nature of their queries.

Usage features - This relates to how each customer uses your service. How frequently they log in, or how long it is since they last logged in, how long they spend on your product, and the actions they take when on there.

Contextual feature - This covers any other data you can gather, like which type of device they access your service on or the customer rep they've contacted the most.

Next, you need to decide the exact question you want your model to answer.

Will the customer churn within a month?

Within three months?

Within six months?

To make it as easy as possible to use your model, create a CSV file and include a row for each customer and a column for each feature.

Step 2 - Create a predictive model

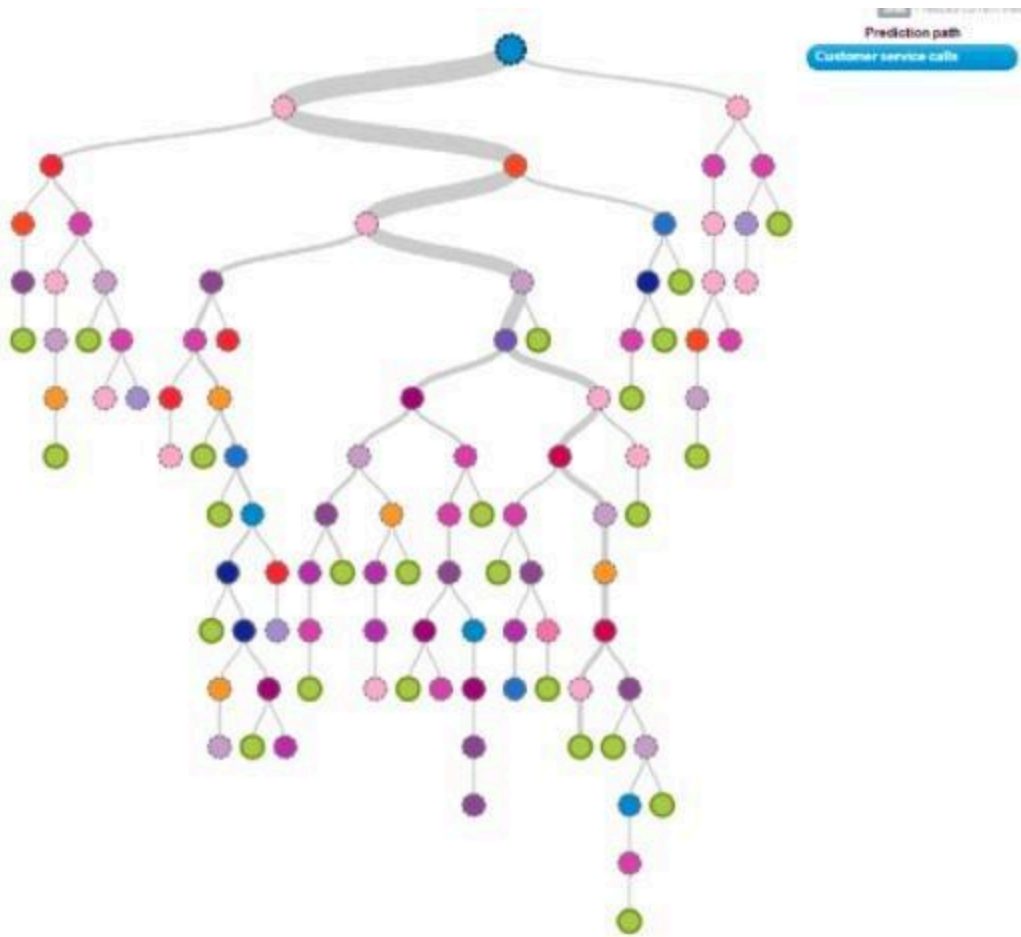
Now you've collected the data, the hardest part is over.

Sit back and let a reputable predictive service do the hard work for you.

Google Cloud ML Engine and **BigML** both create predictive models based on the data that you upload. Just upload the CSV file you've created, and let them do the hard work for you.

We'll use BigML as an example (it's free to create an account), on the dashboard, click on the folder icon and select the file on your computer, or simply drag it from your desktop to the workspace.

What you'll end up with is a decision tree that looks a little something like this...



Each circle represents a question related to your customer’s features, and every branch represents a possible answer. Working your way down the tree, you’ll find a circle with your final output value, stating whether a customer churned or not.

Step 3 - Use the model to predict churn

Now you have a model that understands your churn pattern, you need to use it to predict any future churn.

First of all, you’ll need to collect current data about your existing customers,

This data will match the same data you used to build the model above, except it won’t include the yes or no output for whether or not a customer has churned.

Upload your CSV file to your predictive service, just like we did in the previous step, and generate your predictions!

This will differ slightly depending on the prediction service you choose, but we'll continue with BigML as an example.

Access your model and **select batch prediction** from the lightning bolt icon.

From the **two drop-down menus**, choose your model in the left-hand one and your new customer dataset in the right-hand one.

Select configure, then **output settings**.

Here you can customize how you view your predictions, **select CSV file output** and configure however you like.

Press the **'Predict'** button and on the next page **'Download Batch Prediction'**.

Now you have a file you can open with any spreadsheet program. You'll find your model's predictions in the **'Churn' column**.

Now you have a list of customers at risk of churn, you can focus your energies on retaining them.

P.s. We've got a whole course on [Sales Enablement](#), too.