



## Data Science Case Study

# Predicting B2B Customer Churn

*We built a custom, machine learning model for a leading enterprise software company, helping them identify leading indicators of churn.*



### Industry

Technical Services & Software



### Use Case

B2B Customer Churn



### Techniques

Classification, Adoption & Scale



### Outcome

Accurately & efficiently predicted customers at risk of leaving, giving marketing enough time to combat churn.

## BACKGROUND

A large multinational enterprise software provider noticed a severe drop in maintenance and services renewals specifically in the North American market. The software firm, which supplies a digital business product line across industries to thousands of businesses customers, turned to Mosaic Data Science, a [leading machine learning consulting provider](#)<sup>1</sup>, for help to combat customer churn on services contracts.

To tackle this problem effectively, Mosaic's data science consultants needed to understand the business processes governing the sale of services contracts. Typically, a rep sells a license to a company including maintenance and customer service support for 1 year. A few months before the end of Year 1, the rep sells the company a 1-year maintenance contract for about 20% of the original license fee. A couple months before the end of Year 2, a rep sells the company another 1-year maintenance contract renewal. Ideally, this cycle continues until the software reaches the end of its lifecycle.

Although the initial license fees are profitable, there are upfront costs to the software firm to set up the software and integrate it with a customer's existing IT infrastructure. Maintenance and services fees are much more profitable, so the firm is very motivated to renew these contracts for as many years as possible during a product's life cycle.

If the maintenance sales team knew which customers might not renew, reps could reach out earlier in the year to address customers' concerns and make renewals more likely. Because it would be costly to employ the number of reps it would take to attend to all customers throughout the year, the firm needed a reliable method to identify at-risk clients.

## PREDICTING B2B CUSTOMER CHURN | MODEL SELECTION

Mosaic was asked to develop a Proof of Concept (POC) predictive analytics model to identify leading indicators for customer churn at least 6 months in advance of a contract renewal date. Mosaic's data science consulting team had to collect the correct data. Data sources included maintenance contract agreement dates and renewal dates, customer service call information, and the number of product downloads throughout the contract duration. By subscribing to a maintenance plan, customers are able to freely download any patches or software updates that are released during the contract period. Like service calls, product downloads are an indicator of engagement with the product.

After preparing the data, Mosaic ran a number of different models in R. The data science team decided to predict cancellations with a decision tree algorithm. Compared with other machine learning models, decision trees are prone to overfitting to historical data therefore are often less accurate in their predictions. However, decision trees had two distinct advantages for this project: they are easy to understand, and they provide insight into the most important predictors of the outcome of interest. The interpretability of the models meant that the software firm's



business stakeholders could easily explain the analytic results. Identifying and defining model features that led to churn meant that value would be derived quickly, allowing the maintenance sales team to enhance their at-risk screening process even before the final model went into production. For this customer's needs, these benefits outweighed the risk of lower prediction accuracy.

The Mosaic data science consultant team constructed decision trees to predict cancellations 0, 3, 6, 9, 12, 15, and 18 months in advance of a contract's renewal date. Each of these models took as inputs the number of service calls and downloads over 3-month periods leading up to the prediction point. Figure 1 shows an example of a decision tree for a model predicting cancellations 6 months in advance.

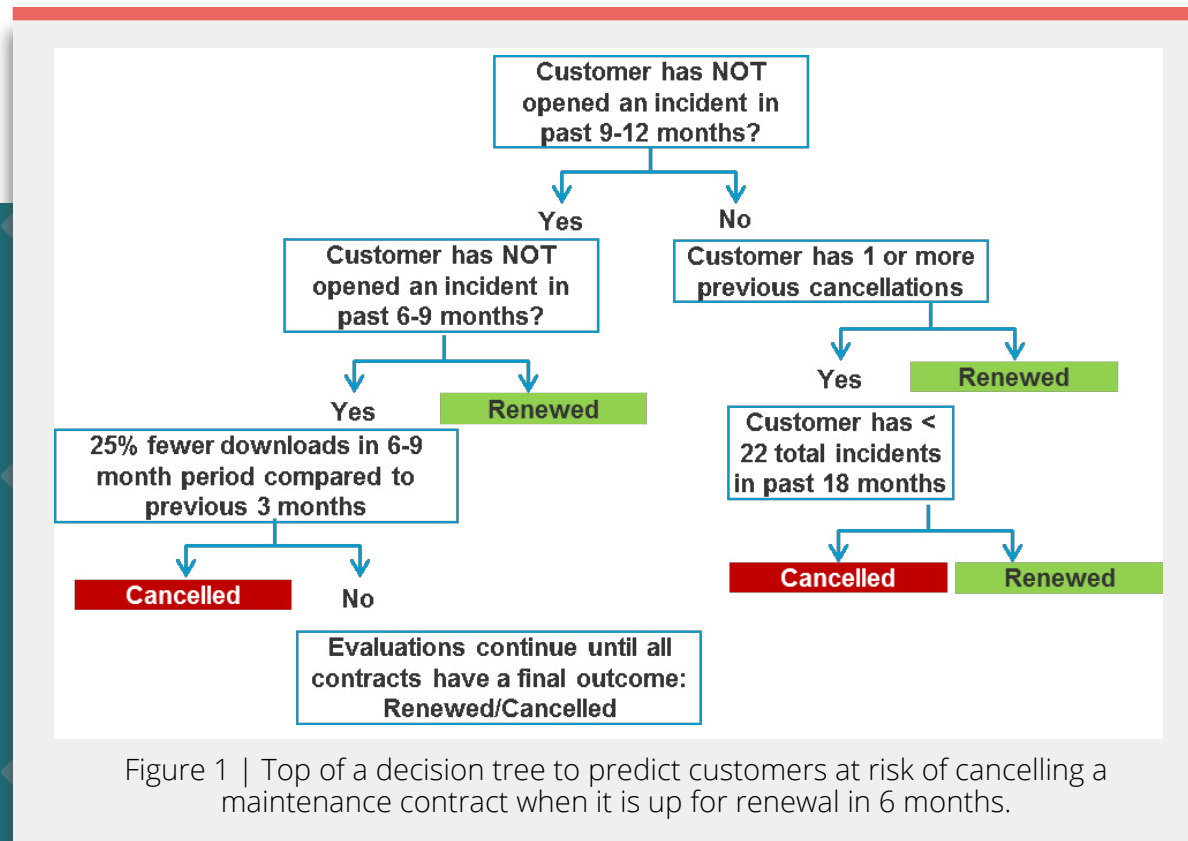


Figure 1 | Top of a decision tree to predict customers at risk of cancelling a maintenance contract when it is up for renewal in 6 months.

## PREDICTING B2B CUSTOMER CHURN | MODEL RESULTS

Because the decision trees developed by the Mosaic team were trained on random samples drawn from the full history of customers, Mosaic measured an average predictive accuracy at each forecast lead time over 1,000 decision trees built on independent samples. The decision trees developed by the Mosaic team are able to predict cancellations with over 70% accuracy on average even 18 months in advance.

To estimate the tradeoff in predictive accuracy due to the selection of decision tree models, the team compared the models' performance against what can be achieved by [support vector machines \(SVMs\)](#)<sup>2</sup>, another popular class of machine learning models. Figure 2 shows the average accuracy of models at each forecast lead time for a decision tree (the black line) and an SVM model (grey line). The greatest performance gap is for predictions made up to 12 months in advance. For 15 and 18 months, the two models perform comparably.

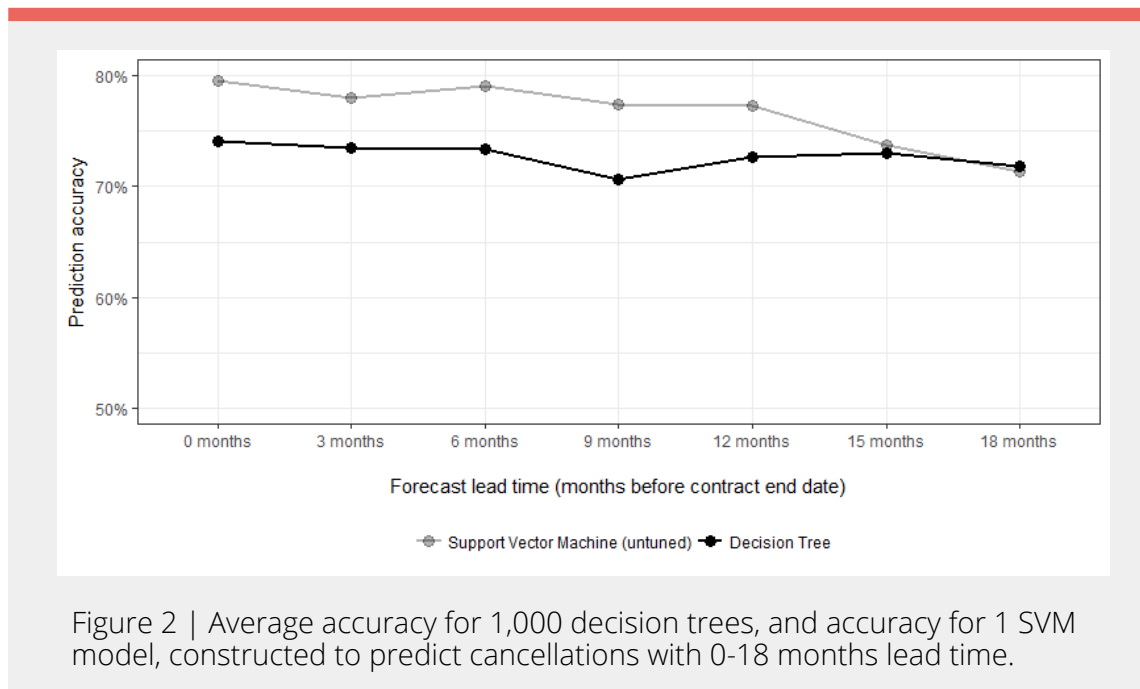


Figure 2 | Average accuracy for 1,000 decision trees, and accuracy for 1 SVM model, constructed to predict cancellations with 0-18 months lead time.



The Mosaic team was able to confirm which features were most impactful to churn. As the software firm expected, the number of product cluster and overall service calls, the volume of service calls about products nearing the end of their maintenance contract, and whether the customer had previously cancelled contracts for other products were all predictors of cancellation of a contract for a particular product. Figure 3 shows the strongest predictors of a cancellation with a 6-month lead time, averaged across 1,000 independently generated decision trees.

Through the application of data science and predictive analytics, the sales team is now able to identify at-risk clients and begin to develop targeted retention practices with the potential to elevate profit margins in the North American market.

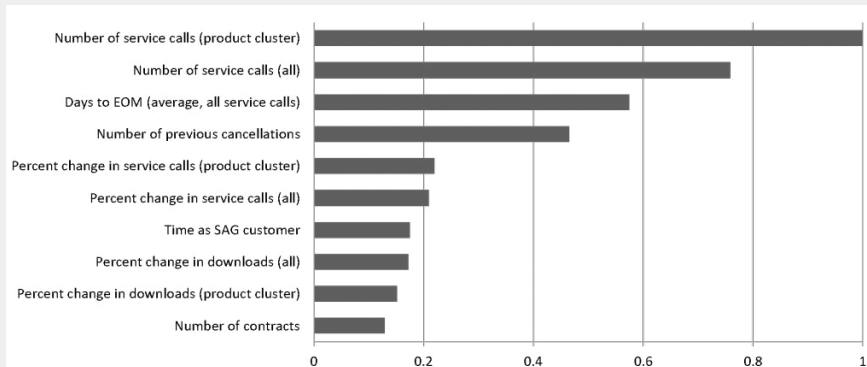


Figure 3 | Average feature importance of predictors of cancellation 6 months in advance. Feature importance measured relative to most important feature (represented by a 1 on the x-axis). EOM: End of Maintenance.

## Endnotes

1. <https://mosaicdatascience.com/tag/professional-services-machine-learning/>
2. [https://en.wikipedia.org/wiki/Support-vector\\_machine](https://en.wikipedia.org/wiki/Support-vector_machine)

