

# Predictive analytics in Azure Machine Learning optimizes credit collections

Every year, Microsoft collects more than \$100 billion in revenue around the world. About 99 percent of financial transactions between customers and Microsoft involve some form of credit. The company's treasury team manages credit and collections for these transactions. The collection process involves all payments—not just late ones—so streamlining and refining a process of this scope is important to our success.

Speeding up collections has a big financial payoff. Considering the amount of revenue, you can safely assume that even small improvements in collection efficiency translate to millions of dollars. And the quicker we collect payments, the quicker we can use that money for activities like extending credit to new customers.

To detect who's likely to pay or not pay—and adjust collection efforts accordingly—Core Services Engineering and Operations (formerly Microsoft IT) partnered with the treasury and finance teams at Microsoft. We brainstormed scenarios, questions, and solutions. We asked things like:

- How do we identify opportunities to improve the collection process?
- How can we predict late payments?
- What technologies and approaches do we use for optimizing credit and collections?
- How do we help the collections team prioritize contacts and decide what actions to take?

## Azure Machine Learning gives us predictive insights

To help with these and other questions, we use data science and Microsoft [Azure Machine Learning](#) as the backbone of our solution. Azure Machine Learning is a cloud-based service that detects patterns in processing large amounts of data, to predict what will happen when you process new data. In other words, it helps us do predictive analytics.

Within two months, we easily set up a predictive model with Azure Machine Learning that helps the collections team prioritize contacts and actions. We use the eXtreme gradient boosting (XGBoost) algorithm—a machine learning method—to create decision trees that answer questions like who's likely to pay versus who isn't.

## How these insights help us in the age of digital transformation

[Data Science for Beginners](#) compares an algorithm to a recipe, and your data to the ingredients. To train and refine the model, we overlay it with five years of historical payment data from our internal database. Output from the model, based on this data, helps us predict with over 80 percent accuracy whether customers are likely to pay late. We also get a valuable understanding of the factors or tendencies linked with customers who've paid versus those who haven't.

Why is this understanding important? The insights we get fit into a broader vision of digital transformation—where we bring together people, data, technology, and processes in new ways to engage customers, empower employees, optimize operations, and transform business solutions.

We use past data and predictive insights from the model to:

- Improve customer satisfaction by reaching out to specific customers with a friendly reminder, while not bothering those who typically pay on time. We prioritize those who've paid late in the past.
- Empower our collections teams, and assign employees to accounts where they're most needed.
- Continuously optimize the efficiency of our collection strategies and business processes.
- Repurpose that money for other short-term and long-term investments.

- Analyze customer behavior and be more predictive and proactive. Beyond deciding which customers to contact first, we see customer trends related to invoice amount, industry, geography, products, and other factors.

## Types of customer insights we get

The insights that we get help us to better understand our markets and to classify customer behavior in those markets. When we onboard new customers, we can correlate certain trends to them quite accurately, based on what we've seen with other customers. Examples include:

- Complex invoices are more likely to be late, and contacting customers with complex invoices by phone helps prevent delays.
- Some customer types and geographies benefit from phone or face-to-face contact much more than others.
- Long-term, high-volume customers and partners are rarely late, and can benefit a lot from payment automation.

## Before we used Azure Machine Learning versus now

Table 1 shows what we used to do, compared to what we do now that we're using Azure Machine Learning, for improving our credit and collections processes.

Table 1. Before versus Now

Before Azure Machine Learning	Now
We often took unnecessary action—for example, contacting customers who aren't likely to pay late. In some ways, it's more about knowing who's likely to pay on time rather than who isn't, so that we avoid contacting those customers.	We mostly contact only customers who need help paying. The benefit is that we can focus on these customers. Also, it provides a good customer experience for those who are likely to pay in any case, because we don't contact them with a reminder.
The collections team contacted every customer with basically the same urgency. The only prioritization was based on balance owed or number of days outstanding. The team first contacted customers who owed the most or who had the most number of days outstanding.	Managers get a list with a risk score that indicates the likelihood that a customer will pay, ordered by the amount that customers owe that month. Managers can then redirect their teams and help prioritize. For example, this person has a 1—they're unlikely to pay on time. Another person has a 0—they're likely to pay on time. So, let's focus on the person with a score of 1.
The collections team used to contact about 90 percent of customers because we lacked the information that we have now.	We need to contact fewer than 40 percent of customers.
We didn't have many insights to speed up how quickly we recovered payments owed or to improve our credit and collections processes. We weren't as predictive or proactive.	We get predictions and insights on areas to improve. We can see trends where customers with certain subscriptions are less likely to pay on time. We know that if customers are in a country/region that's experiencing economic crisis, there's a chance they'll need help paying on time. Contacting them by phone can help us provide solutions faster. Or suppose there's a billing dispute. Based on insights, we correlate that the customer is less likely to pay late because we proactively fix the disputed issue online before the due date.

## Technologies, process, and algorithm that we use

These are the technologies and components that we're using for our solution:

- [Azure Machine Learning Studio](#). It has built-in predictive models, or you can create your own with the R or Python programming languages. We used R and an algorithm called XGBoost.
- Microsoft SQL Server 2014 Enterprise. This is where we store 800 gigabytes of current and historical payment data. From this data, we create categories or features like customer geography, products purchased, purchase frequency, and number of products per order.
- Azure Data Factory. We use this for moving data from SQL Server into Azure Machine Learning, and then bringing the scores back to SQL Server to build reports. These reports contain the invoice information and risk score.
- [Microsoft Power BI](#). We use this to display reports that we use to make collections-related decisions.
- [Microsoft Bot Framework](#) and [Azure App Service](#). We use these for creating our chatbot.

## The solution in action and our iterative process

Azure Machine Learning Studio makes it easy to connect the data to the machine-learning algorithms. Figure 1 below shows the model that we built. We collect data from a variety of data sources and store it in our internal data warehouse called Karnak. Karnak contains historical information from SAP, Microsoft Dynamics CRM Online, MS Sales, our credit-management tool, and external credit bureaus.

We take this data and determine if there are other features that we need to build out of the data to improve the success of the model. This is called feature engineering, and we used this approach to create feature variables such as type of customer, customer tenure, purchase amount, and purchase complexity (products per order).

We then combine the data and engineered features into the machine-learning algorithm called XGBoost to get the late-payment prediction.

Figure 1 quickly summarizes our solution.

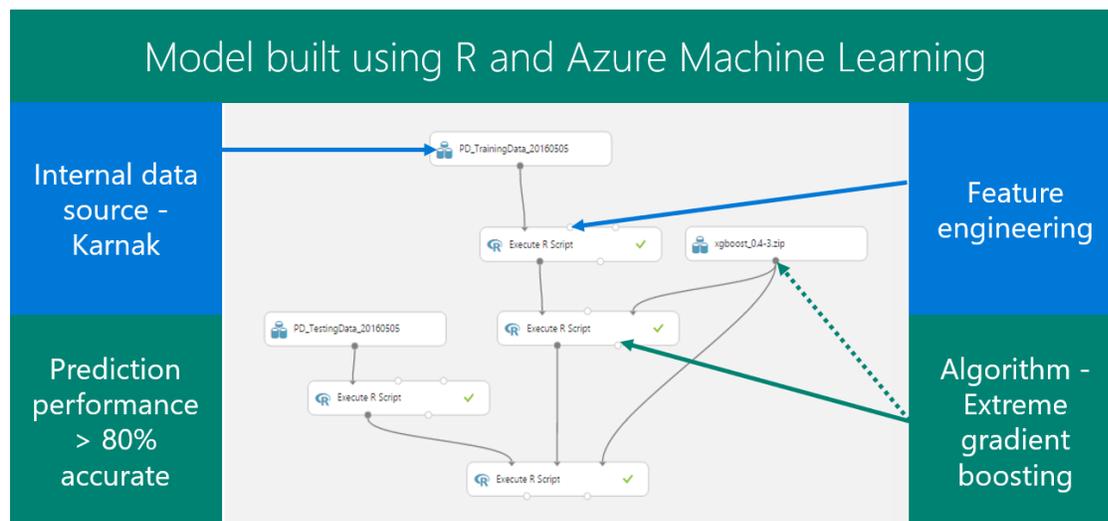


Figure 1. High-level view of the solution

Different skill sets are used within CSEO to build out our machine-learning models. Figure 2 shows the iterative process that we use and the different roles employed at each stage.

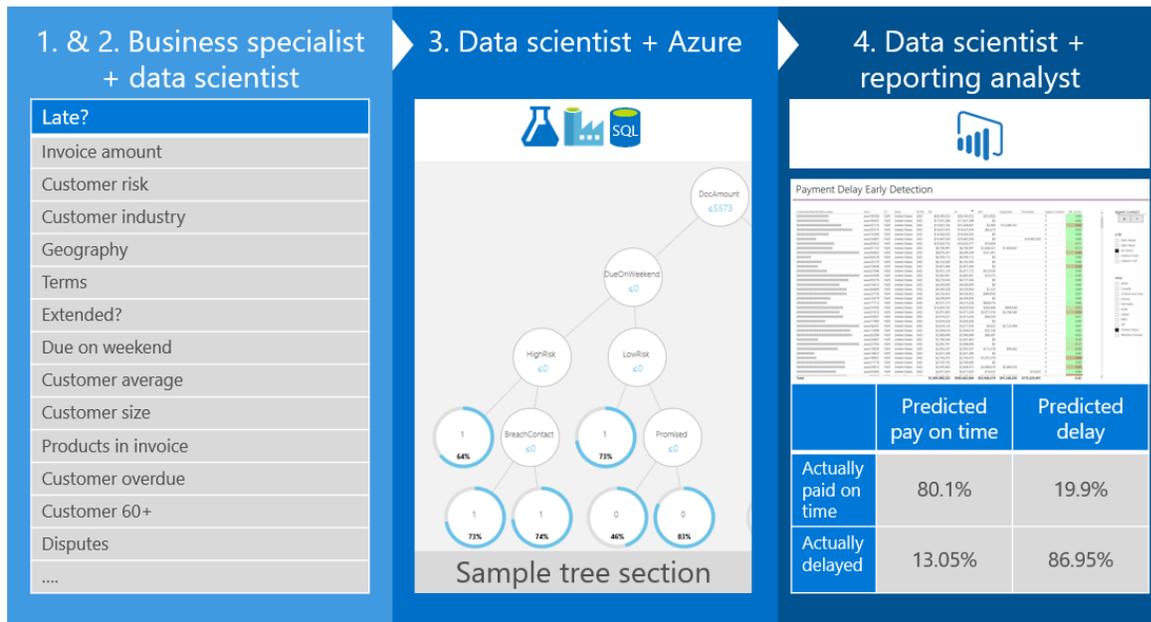


Figure 2. Using Azure Machine Learning for early detection of delayed payments

The process works as follows:

1. The **business specialist** provides information about common business scenarios. In our internal SQL Server data warehouse—Karnak—we have historical data that we separate into features like invoice amount, customer risk, whether a due date has ever been extended, or whether there’s ever been a billing dispute. This data helps predict when a customer will pay. To refine our model, we remove or add features as needed.
2. On an ongoing basis, our **engineers** identify what data we’ll use, and then build a pipeline with data from Karnak to enable the predictive model. They decide how the data should go into the model, how to retrieve the score, and where the data from the model will be stored. Then they provide the data to the **data science team**.
3. Next, the **data scientists** create the model by using Azure Machine Learning Studio, Azure SQL Database, and Azure Data Factory. Data is put into decision trees.

*Note: The decision tree in Figure 2 is for illustrative purposes only. We have more than 1,000 trees. The largest tree has 100 levels.*

This produces one of two outcomes:

- Late = 0 (a customer pays on time)
- Late = 1 (a customer pays late)

Azure Machine Learning also gives us a risk percentage score of how likely the customer is to pay on time.

4. The scores go into our Karnak database and are displayed in Power BI reports to collections teams.

## Our chatbot helps us quickly answer recurring questions

Credit and collections team members often come across the same questions over and over. To speed up the process of answering these recurring questions, we built a chatbot. There are thousands of questions in emails, but there wasn't a real tracking system.

The following steps, as shown in Figure 3, show how the chatbot works:

1. The user asks a question to the chatbot in plain English.
2. The chatbot uses Language Understanding Service (LUIS) to translate the question from plain English to a computer-understandable language.
3. The chatbot asks a question to a web service that connects to Karnak, our internal credit-data mall.
4. The chatbot formats and presents an answer to the user.

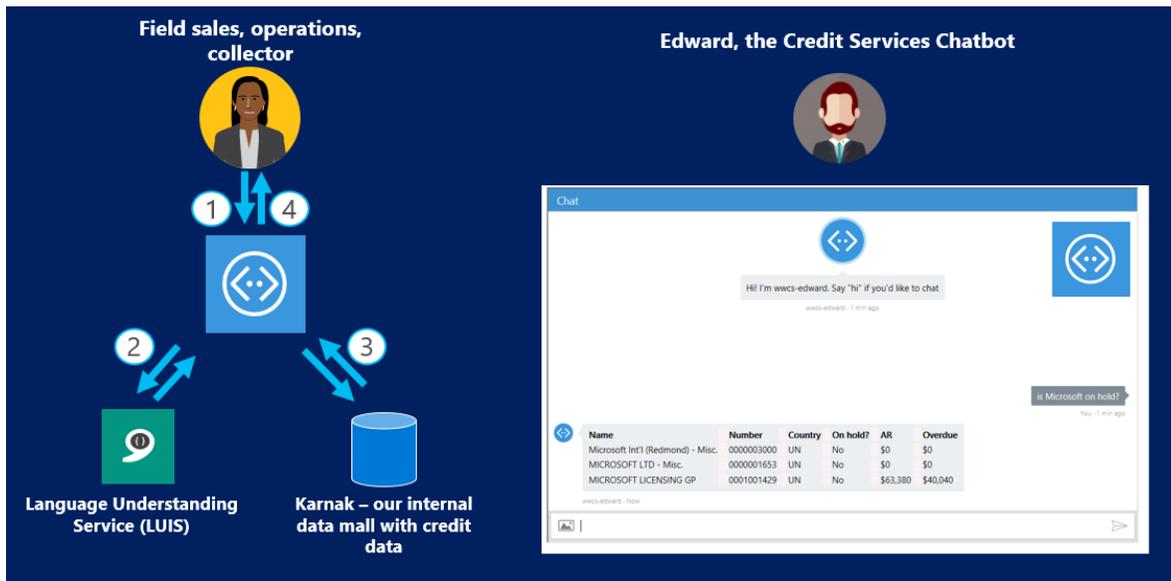


Figure 3. The chatbot that we created

Now, field sales, operations, and collectors can see the latest information about customers they interact with and detect issues. For example, they easily see what the customer credit limit is, the overdue amount, whether a customer has exceeded the credit limit and is temporarily blocked, and answers to other questions.

We used Bot Framework and Azure App Service. The chatbot talks to App Service, and App Service talks to Karnak. Karnak data goes into Azure SQL Database, and App Service connects to SQL Database to answer the bot's questions.

## The eXtreme gradient boosting algorithm helps us prioritize

We use the XGBoost algorithm to create decision trees that look at features. For example, suppose an invoice is due on Saturday, or a customer in a particular country/region tends to pay late, and the average invoice is, say, \$2,000. After we have the forest of trees that explain the historical data, we put new data in different trees. If most of the trees predict that an invoice will be late, we mark it accordingly. Otherwise, we mark it as unlikely to be late.

For customers with invoices that are due soon, the model shows which customers to prioritize. It puts their names at the top of a list for the collectors, so that they can contact these customers earlier in the process.

## Challenges we faced

Here are some of the challenges that we initially had, but that we overcame:

- **Eliminating security-related fears.** We proved that our architecture is secure and robust. But even in IT organizations like ours, there can be people who worry about putting sensitive data in the cloud. That fear is now far less common at Microsoft, but it's possible that IT departments at some companies might needlessly worry. In our case, Azure was protecting high business impact (HBI) data, and the security auditing team was confident in Azure security even for mission-critical finance information. Azure offers strong levels of security that meet our compliance standards for HBI finance data. Our internal risk team reviews all applications, and all of these applications are tested for data and system compliance.
- **Finding people who understand the business, data science, and machine learning, and who are effective liaisons.** Finding people with deep knowledge of these areas takes time in itself. But it's also crucial that they bridge communication with the business data science team. After that, setting up a predictive model is quick.
- **Convincing everyone how well the predictive model works.** We proved that the model wasn't guessing randomly and that the prediction performance has a high degree of accuracy. Proving this early on can make it quicker to move the solution to production.

## Lessons learned

To have the right data to put into an algorithm, you should have someone who understands the business processes and has good business insights. In our case, we had people with this knowledge and five years of historical data. We knew what business factors were important. But say you're starting from scratch. If you don't have someone who understands the business scenarios, and you don't have much historical data, it's harder.

Solving the machine learning problem itself took us only about two months, but deploying it took longer. There were lots of reviews and test cycles to demonstrate the accuracy and the high level of security that we have. If you're doing something similar, build in extra time to allow for these cycles.

It's unreasonable to assume you'll get it perfect the first time. To get expected, consistent results, keep iterating.

## Looking ahead

We keep learning all the time as we iterate. Down the road, we plan to build on what we're doing now. For example, we have integrated insights into several of our collection processes and some systems, but not all of them. Our approach is to incorporate changes to get the best return, and we're still working on deploying these AI-based insights to everything we do.

We have also started to expand our scenarios into areas that are adjacent to credit and collections: sales and supply-chain features. We plan to add additional scenarios, use cases, data sources, and data-science resources for even more insights.

Also on our feature list is macroeconomic data, such as gross domestic product, inflation, and foreign exchange, to make our predictions even better.

## Conclusion

Even small improvements in collections efficiency add up to millions of dollars. In an age of digital transformation, data and predictive insights are key assets that help us tailor our strategies and focus our efforts on what's most important. With data science, Azure Machine Learning, and predictive analytics, we improve customer satisfaction, empower our collections team, optimize the efficiency and speed of our collection operations, and we're more predictive and proactive.

## For more information

### Microsoft IT Showcase

[microsoft.com/ITShowcase](https://microsoft.com/ITShowcase)

[Azure Machine Learning Studio](#)

[IT Showcase Data and AI content](#)

[Transforming business through the power of data](#)

© 2018 Microsoft Corporation. This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY. The names of actual companies and products mentioned herein may be the trademarks of their respective owners.