Improving Customer Support in Electronics Manufacturing

Understanding and getting ahead of product support issues

About the Customer

The company is one of the leading manufacturers of electronic products for the consumer and professional markets. Its diversified business includes consumer and professional electronics, gaming, entertainment, and financial services.

This case study focuses on its electronics manufacturing business, specifically in Europe. The electronics business is organized into largely independent geographic divisions, with groups in Japan, China, Asia Pacific (other than Japan and China), the Americas, and Europe, which includes Russia and Turkey.

The RapidMiner user group within the European electronics business is a data science team that is part of the organization focused on post-sales customer support.

The Customer's Need

Decrease cost of customer support

As a manufacturer of high-end consumer electronics, the company must, and does, offer post-sales support to its customers. This includes giving customers the opportunity to call or email the company to ask questions about the products they have purchased, including requesting technical assistance.

The company's objective is always to reduce the cost of providing this support, while at the same time ensuring customers have a positive experience with their products. Strategies to achieve this can vary from improving the content on the website, so customers can self-serve instead of calling for help, to better preparing call center agents for quick resolution of customer calls, to building products that require less support to begin with.

The data science team in the company's post-sales organization has the mission of using data to help achieve these goals.

The Challenge

Before RapidMiner: analysis of "what" not "why"

Several years ago, before this manufacturer's data science team adopted RapidMiner, it did a good job of understanding the basic statistics as they related to customer support - such as how many people called, how long people stayed on the phone, how many people visited the support website, etc. But it was more difficult for the team to determine why people were calling, or whether the customer support team was actually able to provide solutions for customers when they called.

The data science team primarily used structured data in conducting its analyses, because this data was easier to work with. But the data had its limitations, as it only allowed the kinds of analyses mentioned - basic statistical summaries, no determination of causality, or other deeper insights.

They wanted to do more than this, and tried to, on occasion. For example, when the team would learn about a particular issue that customers had been calling in about frequently, they could dig into their unstructured data, such as call logs and emails from customers, and do queries on specific words or terms to look for insights. But these only helped the business better understand the issues that had already been identified as a problem. In many cases, it was much too late to do anything about the problem in a way that could create significant savings for the business. In most cases, this additional analysis was merely confirming things the business already knew.

The team used Business Objects as its primary tool at this time, as these types of summarizations is what a BI (business intelligence) tool is designed for. The team had rudimentary approaches and tools to help with other aspects of the work - such as manual usage of Google Translate on those occasions when they tried to dig into the call logs.

The Solution

RapidMiner supports the first steps towards leveraging unstructured data

The data science team knew they could do more to support the business goal of reducing support costs if they could make better use of the vast quantities of unstructured data they had available to them. Even with its regional focus on Europe, the data represented an enormous volume of customer interactions, covering calls from Spain to Siberia.

The data science team proactively sought ways to make better use of this data. While the mandate from company headquarters is always “do more with less,” the data science team wasn’t told specifically to dig into unstructured data. But they knew there was great potential in unlocking that data, it could dramatically impact the mission of the post-sales organization. The more they could learn about “why” people were contacting them, rather than just how many and often they were contacted, the more of an impact they could have.

Further, the team knew to cast a wide net when seeking out sources of unstructured data. Rather than limiting themselves to the (relatively) easily-obtainable call logs and email inquiries from customers, they collected online product reviews, social media comments, and as many other sources as they could. This allowed them to analyze the people who didn’t contact the company, a group just as important as the people who did contact them. Similarly, the team set out to capture not just what people said about the company itself, but also what they said about competitor’s products, as a comparison. This was important not just to have a variety of sources, but because the products are becoming more complex, the data science team needs as much information as possible to cover every possible aspect of product usage and need for support.

A first step, in the direction of deeper analytical insight and greater business value, was to focus on classification analyses - why people are calling, and in as much detail as possible (reasons and multiple layers of sub-reasons). To do this, the team first had to automate many of its existing business processes. It couldn’t simply hire hundreds of people to read customer comments.

This was how the team first started using RapidMiner, they were able to quickly improve data preparation and other ETL-type work. And initially, that was the extent of their usage of the RapidMiner platform, the team was essentially automating work people were already doing. Still, this was important first step.

An example of this was the translations process. With RapidMiner, the team could easily create workflows that allowed its unstructured data, in 26 different languages, to be routinely translated for easier interaction and analysis. The power of the RapidMiner platform was already shining, even in these simple, initial applications.

RapidMiner helps bring machine learning to bear

The next step was to bring machine learning to bear on the problem. At the time, several years ago, machine learning was a new topic in the industry, and at this particular company as well, with lots of experimentation but not a lot of practical examples or guidance to draw on. So, the team joined other industry pioneers in the use of machine learning to support the business needs of manufacturers. And in doing so, they began to harness the real power of the RapidMiner platform.

This electronics manufacturer first used RapidMiner's machine learning capabilities for web and text mining, to support its classification analyses. RapidMiner allows the team to understand more deeply what's going on with customer support calls. They can identify trends, such as peaks and valleys of calls, and more importantly the reasons behind those calls - what problems people are having, and whether they're finding resolutions. By analyzing the unstructured data, and not just limiting themselves to the structured data about customer calls, they can get to the "why," not just the "what."

Now the team is moving on to do even more powerful analysis with RapidMiner and machine learning, such as:

- **Anomaly detection:** to easily tell the difference between routine support inquiries and signs that more inbound inquiries will be coming on a specific topic
- **Long-tail problem identification:** to help find lots of smaller problems, that might easily be missed, but that if addressed could create a big impact for the business, rather than just focusing on the big, obvious problems
- **Forecasting: to answer questions,** such as - Are there enough repair parts in stock to be prepared for the expected volume of service calls?

These predictive analyses are an entirely new approach to the problem for the company. Before RapidMiner, its data science team occasionally did intensive logistic regressions. But for the most part, the company started to almost entirely look forward, rather than backwards, when it began using RapidMiner to the fullest.

RapidMiner's insights drive real business changes

What does the company do with these insights from the data science team? It uses them to adjust practically every aspect of its operations to reduce customer support costs and improve its customer experience.

For example, the company can arm its call center team with the information needed to handle expected calls from customers quickly. Its customer support team can refine the content on its web site to make sure it covers topics customers are likely to struggle with, so customers can access the content and not need to call in or email. The marketing team can adjust its communication strategy to adjust customer perception of the products along the lines identified in the analysis. And the company can even feed these insights back to the product development and manufacturing organizations, so they are better able to build products people won't need to support with at all.

RapidMiner is at the core of the company's machine learning technical ecosystem

RapidMiner is an integral part of a larger technical ecosystem that the company has built to support its data science. They have connected RapidMiner to multiple databases, a search engine, Python libraries (called and controlled by RapidMiner), and more, to accomplish its goals. Business Objects continues to be a frequently-used tool, as well. But RapidMiner is the core of the entire ecosystem.

The Results

Analytics efficiency and precision now, business KPIs next

For now, the data science team is measuring its success based on how much more efficiently it can produce the kind of insights that drive business success. Previously, the data science team operated in a kind of 80/20 environment - 80% of its time was spent collecting and managing data, and only 20% analyzing it. With RapidMiner, the company is on its way to flipping this on its head, with 80% of time being spend on analyses that drive business value. In addition, the team always tracks the precision and accuracy of its predictions, which is strong and getting even stronger.

This efficiency enables the team to do new things that would have been impossible before. For example, for social network analysis, the data science team uses RapidMiner to crawl websites on a daily basis, harvesting and analyzing reviews and sharing the results with business users. Without RapidMiner, the company would never have tried to do this, or might have, but on a much smaller sample set - and with much less business impact.

However, the team also plans to start to track the bottom-line impact of its analytical insights, using the same business KPIs as its partner organizations do - such as customer support, marketing, and product development. This may ultimately be the best indicator of the impact of the team's new approach to analytics with RapidMiner: that it naturally brings the data science team extremely close to the business results they help drive.

